

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification : Not classified		A2	(11) International Publication Number: WO 97/03549 (43) International Publication Date: 6 February 1997 (06.02.97)
(21) International Application Number: PCT/US96/11920 (22) International Filing Date: 18 July 1996 (18.07.96) (30) Priority Data: 60/001,498 19 July 1995 (19.07.95) US (71) Applicant (for all designated States except US): ASCOM NEXION INC. [US/US]; 289 Great Road, Acton, MA 01720 (US). (72) Inventors; and (75) Inventors/Applicants (for US only): HAUSER, Stephen, A. [US/US]; 106 Farms Drive, Burlington, MA 01803 (US). CALDARA, Stephen, A. [US/US]; 220 Horsepond Road, Sudbury, MA 01776 (US). MANNING, Thomas, A. [US/US]; 26 Summer Street, Northboro, MA 01532 (US). (74) Agents: LEBOVICI, Victor, B. et al.; Weingarten, Schurgin, Gagnebin & Hayes, Ten Post Office Square, Boston, MA 02109 (US).			(81) Designated States: AL, AM, AT, AU, AZ, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TR, TT, UA, UG, US, UZ, VN, ARIPO patent (KE, LS, MW, SD, SZ, UG), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG). Published Without international search report and to be republished upon receipt of that report.
(54) Title: PRIORITIZED ACCESS TO SHARED BUFFERS			
(57) Abstract <p>In a link-level flow controlled system, a method and apparatus providing the ability to partition a buffer resource among multiple prioritized buffer subsets through definition of at least one threshold, the buffer resource being shared by a plurality of connections. Different categories of service levels, in terms of delay bounds, are thus enabled. The presently disclosed link-level flow controlled system provides for zero cell loss. The shared buffer resource is divided among N priority pools, defined by N-1 threshold levels, each priority pool attributable to a respective category of service. Link-level counters and registers, disposed in a transmit element, as well as an indication of priority level associated with each connection, are employed in realizing the shared buffer resource.</p>			

AL

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

- 1 -

PRIORITIZED ACCESS TO SHARED BUFFERS

5

FIELD OF THE INVENTION

10

This application relates to communications methods and apparatus in a distributed switching architecture, and in particular to methods and apparatus for prioritized access to shared buffer resources in a flow controlled virtual connection.

15

BACKGROUND OF THE INVENTION

20

A Flow Controlled Virtual Connection (FCVC) protocol for use in a distributed switching architecture is presently known in the art, and is discussed below with reference to Fig. 1. This protocol involves communication of status (buffer allocation and current state) on a per virtual connection, such as a virtual channel connection or a virtual path connection, basis between upstream and downstream network elements to provide a "no cell loss" guarantee. A cell is the unit of data to be transmitted. Each cell

25

requires a buffer to store it. Presently known flow control mechanisms further provide for the application of connection-level flow control among a plurality of connections, each having equivalent access to a shared buffer pool in the receiver. No distinction in buffer allocation is made within a service category, regardless of time-sensitivity of one or more of the data cell sequences.

30

SUMMARY OF THE INVENTION

35

In a link-level flow controlled system, the ability to partition a buffer resource among multiple prioritized buffer

- 2 -

subsets through definition of at least one threshold is provided, the buffer resource for sharing by a plurality of connections. Different categories of service levels in terms of delay bounds are thus enabled. The link-level flow controlled system according to the present invention provides for zero cell loss by the throttling of cell flow from a transmitter element to a receiver element based upon link-level counters and registers maintained within the transmitter element updated by feedback information from the receiver element. A link can be defined either as a physical link or as a logical grouping comprised of logical connections.

Access to the partitioned buffer subsets is based upon the priority assigned to each subset and a priority associated with each connection having cells to be buffered. Buffer subsets are defined by the at least one threshold, which in one embodiment is dynamically adjustable.

Link-level counters and registers found in the transmitter element maintain threshold levels and track buffer pool usage in the receiver element for controlling cell traffic from the transmitter element to the receiver element. As with link-level flow control, feedback is provided from the receiver element reflective of the current status of the shared buffer resource.

BRIEF DESCRIPTION OF THE DRAWINGS

The above and further advantages may be more fully understood by referring to the following description and accompanying drawings of which:

Fig. 1 is a block diagram of a connection-level flow control apparatus as known in the prior art;

Fig. 2 is a block diagram of a link-level flow control apparatus according to the present invention;

Figs. 3A and 3B are flow diagram representations of counter initialization and preparation for cell transmission within a flow control method according to the present

- 3 -

invention;

Fig. 4 is a flow diagram representation of cell transmission within the flow control method according to the present invention;

5 Figs. 5A and 5B are flow diagram representations of update cell preparation and transmission within the flow control method according to the present invention;

10 Figs. 6A and 6B are flow diagram representations of an alternative embodiment of the update cell preparation and transmission of Figs. 5A and 5B;

Figs. 7A and 7B are flow diagram representations of update cell reception within the flow control method according to the present invention;

15 Figs. 8A, 8B and 8C are flow diagram representations of check cell preparation, transmission and reception within the flow control method according to the present invention;

20 Figs. 9A, 9B and 9C are flow diagram representations of an alternative embodiment of the check cell preparation, transmission and reception of Figs. 8A, 8B and 8C;

Fig. 10 illustrates a cell buffer pool according to the present invention as viewed from an upstream element;

25 Fig. 11 is a block diagram of a link-level flow control apparatus in an upstream element providing prioritized access to a shared buffer resource in a downstream element according to the present invention;

Figs. 12A and 12B are flow diagram representations of counter initialization and preparation for cell transmission within a prioritized access method according to the present invention;

30 Figs. 13A and 13B illustrate alternative embodiments of cell buffer pools according to the present invention as viewed from an upstream element;

35 Fig. 14 is a block diagram of a flow control apparatus in an upstream element providing guaranteed minimum bandwidth and prioritized access to a shared buffer resource in a downstream element according to the present invention;

- 4 -

Figs. 15A and 15B are flow diagram representations of counter initialization and preparation for cell transmission within a guaranteed minimum bandwidth mechanism employing prioritized access according to the present invention;

Fig. 16 is a block diagram representation of a transmitter, a data link, and a receiver in which the presently disclosed joint flow control mechanism is implemented; and

Fig. 17 illustrates data structures associated with queues in the receiver of Fig. 16.

DETAILED DESCRIPTION

In Fig. 1, the resources required for connection-level flow control are presented. As previously stated, the illustrated configuration of Fig. 1 is presently known in the art. However, a brief discussion of a connection-level flow control arrangement will facilitate an explanation of the presently disclosed link-level flow control method and apparatus.

One link 10 is shown providing an interface between an upstream transmitter element 12, also known as an UP subsystem, and a downstream receiver element 14, also known as a DP subsystem. Each element 12, 14 can act as a switch between other network elements. For instance, the upstream element 12 in Fig. 1 can receive data from a PC (not shown). This data is communicated through the link 10 to the downstream element 14, which in turn can forward the data to a device such as a printer (not shown). Alternatively, the illustrated network elements 12, 14 can themselves be network end-nodes.

The essential function of the presently described arrangement is the transfer of data cells from the upstream element 12 via a connection 20 in the link 10 to the downstream element 14, where the data cells are temporarily held in cell buffers 28. Cell format is known, and is further described in "Quantum Flow Control", Version 1.5.1,

- 5 -

dated June 27, 1995 and subsequently published in a later version by the Flow Control Consortium. In Fig. 1, the block labelled Cell Buffers 28 represents a set of cell buffers dedicated to the respective connection 20. Data cells are released from the buffers 28, either through forwarding to another link beyond the downstream element 14, or through cell utilization within the downstream element 14. The latter event can include the construction of data frames from the individual data cells if the downstream element 14 is an end-node such as a work station.

Each of the upstream and downstream elements 12, 14 are controlled by respective processors, labelled UP (Upstream Processor) 16 and DP (Downstream Processor) 18. Associated with each of the processors 16, 18 are sets of buffer counters for implementing the connection-level flow control. These buffer counters are each implemented as an increasing counter/limit register set to facilitate resource usage changes. The counters of Fig. 1, described in further detail below, are implemented in a first embodiment in UP internal RAM. The counter names discussed and illustrated for the prior art utilize some of the same counter names as used with respect to the presently disclosed flow control method and apparatus. This is merely to indicate the presence of a similar function or element in the prior art with respect to counters, registers, or like elements now disclosed.

Within the link 10, which in a first embodiment is a copper conductor, multiple virtual connections 20 are provided. In an alternative embodiment, the link 10 is a logical grouping of plural virtual connections 20. The number of connections 20 implemented within the link 10 depends upon the needs of the respective network elements 12, 14, as well as the required bandwidth per connection. In Fig. 1, only one connection 20 and associated counters are illustrated for simplicity.

First, with respect to the upstream element 12 of Fig. 1, two buffer state controls are provided, BS_Counter 22 and

- 6 -

BS_Limit 24. In a first embodiment, each are implemented as fourteen bit counters/registers, allowing a connection to have 16,383 buffers. This number would support, for example, 139 Mbps, 10,000 kilometer round-trip service. The buffer state counters 22, 24 are employed only if the connection 20 in question is flow-control enabled. That is, a bit in a respective connection descriptor, or queue descriptor, of the UP 16 is set indicating the connection 20 is flow-control enabled.

BS_Counter 22 is incremented by the UP 16 each time a data cell is transferred out of the upstream element 12 and through the associated connection 20. Periodically, as described below, this counter 22 is adjusted during an update event based upon information received from the downstream element 14. BS_Counter 22 thus presents an indication of the number of data cells either currently being transmitted in the connection 20 between the upstream and downstream elements 12, 14, or yet unreleased from buffers 28 in the downstream element 14.

BS_Limit 24 is set at connection configuration time to reflect the number of buffers 28 available within the receiver 14 for this connection 20. For instance, if BS_Counter 22 for this connection 20 indicates that twenty data cells have been transmitted and BS_Limit 24 indicates that this connection 20 is limited to twenty receiver buffers 28, the UP 16 will inhibit further transmission from the upstream element 12 until an indication is received from the downstream element 14 that further buffer space 28 is available for that connection 20.

Tx_Counter 26 is used to count the total number of data cells transmitted by the UP 16 through this connection 20. In the first embodiment, this is a twenty-eight bit counter which rolls over at 0xFFFFFFFF. As described later, Tx_Counter 16 is used during a check event to account for errored cells for this connection 20.

In the downstream element 14, the DP 18 also manages a

- 7 -

set of counters for each connection 20. Buffer_Limit 30 performs a policing function in the downstream element 14 to protect against misbehaving transmitters. Specifically, the buffer_limit register 30 indicates the maximum number of cell buffers 28 in the receiver 14 which this connection 20 can use. In most cases, BS_Limit 24 is equal to Buffer_Limit 30. At some point, though, it may be necessary to adjust the maximum number of cell buffers 28 for this connection 20 up or down. This function is coordinated by network management software. To avoid the "dropping" of data cells in transmission, an increase in buffers per connection is reflected first in Buffer_Limit 30 prior to BS_Limit 24. Conversely, a reduction in the number of receiver buffers per connection is reflected first in BS_Limit 24 and thereafter in Buffer_Limit 30.

Buffer_Counter 32 provides an indication of the number of buffers 28 in the downstream element 14 which are currently being used for the storage of data cells. As described subsequently, this value is used in providing the upstream element 12 with a more accurate picture of buffer availability in the downstream element 14. Both the Buffer_Limit 30 and Buffer_Counter 32 are fourteen bits wide in the first embodiment.

N2_Limit 34 determines the frequency of connection flow-rate communication to the upstream transmitter 12. A cell containing such flow-rate information is sent upstream every time the receiver element 14 forwards a number of cells equal to N2_Limit 34 out of the receiver element 14. This updating activity is further described subsequently. In the first embodiment, N2_Limit 34 is six bits wide.

The DP 18 uses N2_Counter 36 to keep track of the number of cells which have been forwarded out of the receiver element 14 since the last time the N2_Limit 34 was reached. In the first embodiment, N2_Counter 36 is six bits wide.

In a first embodiment, the DP 18 maintains Fwd_Counter 38 to maintain a running count of the total number of cells

- 8 -

forwarded through the receiver element 14. This includes buffers released when data cells are utilized for data frame construction in an end-node. When the maximum count for this counter 38 is reached, the counter rolls over to zero and continues. The total number of cells received by the receiver element 14 can be derived by adding Buffer_Counter 32 to Fwd_Counter 38. The latter is employed in correcting the transmitter element 12 for errored cells during the check event, as described below. Fwd_Counter 38 is twenty-eight bits wide in the first embodiment.

In a second embodiment, the DP 18 maintains Rx_Counter 40, a counter which is incremented each time the downstream element 14 receives a data cell through the respective connection 20. The value of this counter 40 is then usable directly in response to check cells and in the generation of an update cell, both of which will be described further below. Similar to the Fwd_Counter 38, Rx_Counter 40 is twenty-eight bits wide in this second embodiment.

There are two events in addition to a steady state condition in the connection-level flow controlled protocol: update; and check. In steady state, data cells are transmitted from the transmitter element 12 to the receiver element 14. In update, buffer occupancy information is returned upstream by the receiver element 14 to correct counter values in the transmitter element 12. Check mode is used to check for cells lost or injected due to transmission errors between the upstream transmitter and downstream receiver elements 12, 14.

In the accompanying figures, connection level counters are augmented with "[i]" to indicate association with one connection [i] of plural possible connections.

Prior to any activity, counters in the upstream and downstream elements 12, 14 are initialized, as illustrated in Fig. 3A. Initialization includes zeroing counters, and providing initial values to limit registers such as Link_BS_Limit and Link_Buffer_Limit. In Fig. 3A,

- 9 -

Buffer_Limit[i] is shown being initialized to $(RTT \cdot BW) + N2$, which represents the round-trip time times the virtual connection bandwidth, plus accommodation for delays in processing the update cell. As for Link_N2_Limit, "X" represents the buffer state update frequency for the link, and for N2_Limit[i], "Y" represents the buffer state update frequency for each connection.

In steady state operation, the UP 16 of the transmitter element 12 determines which virtual connection 20 (VC) has a non-zero cell count (i.e. has a cell ready to transmit), a BS_Counter value less than the BS_Limit, and an indication that the VC is next to send (also in Figs. 3A and 3B).

The UP 16 increments BS_Counter 22 and Tx_Counter 26 whenever the UP 16 transmits a data cell over the respective connection 20, assuming flow control is enabled (Fig. 4). Upon receipt of the data cell, the DP 18 checks whether Buffer_Counter 32 equals or exceeds Buffer_Limit 30, which would be an indication that there are no buffers available for receipt of the data cell. If Buffer_Counter \geq Buffer_Limit, the data cell is discarded (Fig. 3B). Otherwise, the DP 18 increments Buffer_Counter 32 and Rx_Counter 40 and the data cell is deposited in a buffer cell 28, as in Fig. 4. The Tx_Counter 26 and the Rx_Counter 40 roll over when they reach their maximum.

If flow control is not enabled, none of the presently described functionality is implemented. Connections that do not utilize flow control on the link can coexist with connections using link flow control. The flow control accounting is not employed when cells from non-flow controlled connections are transmitted and received. This includes both connection level accounting and link level accounting. Thereby, flow control and non-flow control connections can be active simultaneously.

When a data cell is forwarded out of the receiver element 14, Buffer_Counter 32 is decremented. Buffer_Counter 32 should never exceed Buffer_Limit 30 when the connection-

- 10 -

level flow control protocol is enabled, with the exception of when BS_Limit 24 has been decreased and the receiver element 14 has yet to forward sufficient cells to bring Buffer_Counter 32 below Buffer_Limit 30.

5 A buffer state update occurs when the receiver element 14 has forwarded a number of data cells equal to N2_Limit 34 out of the receiver element 14. In the first embodiment in which the DP 18 maintains Fwd_Counter 38, update involves the transfer of the value of Fwd_Counter 38 from the receiver
10 element 14 back to the transmitter element 12 in an update cell, as in Fig. 6A. In the embodiment employing Rx_Counter 40 in the downstream element 14, the value of Rx_Counter 40 minus Buffer_Counter 32 is conveyed in the update cell, as in Fig. 5A. At the transmitter 12, the update cell is used
15 to update the value in BS_Counter 22, as shown for the two embodiments in Fig. 7A. Since BS_Counter 22 is independent of buffer allocation information, buffer allocation can be changed without impacting the performance of this aspect of connection-level flow control.

20 Update cells require an allocated bandwidth to ensure a bounded delay. This delay needs to be accounted for, as a component of round-trip time, to determine the buffer allocation for the respective connection.

25 The amount of bandwidth allocated to the update cells is a function of a counter, Max_Update_Counter (not illustrated) at an associated downstream transmitter element (not illustrated). This counter forces the scheduling of update and check cells, the latter to be discussed subsequently. There is a corresponding Min_Update_Interval
30 counter (not shown) in the downstream transmitter element, which controls the space between update cells. Normal cell packing is seven records per cell, and Min_Update_Interval is similarly set to seven. Since the UP 16 can only process one update record per cell time, back-to-back, fully packed
35 update cells received at the UP 16 would cause some records to be dropped.

- 11 -

An update event occurs as follows, with regard to Figs. 1, 5A and 6A. When the downstream element 14 forwards (releases) a cell, Buffer_Counter 32 is decremented and N2_Counter 36 and Fwd_Counter 38 are incremented. When the N2_Counter 36 is equal to N2_Limit 34, the DP 18 prepares an update cell for transmission back to the upstream element 12 and N2_Counter 36 is set to zero. The upstream element 12 receives a connection indicator from the downstream element 14 forwarded cell to identify which connection 20 is to be updated. In the first embodiment, the DP 18 causes the Fwd_Counter 38 value to be inserted into an update record payload (Fig. 6A). In the second embodiment, the DP 18 causes the Rx_Counter 40 value minus the Buffer_Counter 32 value to be inserted into the update record payload (Fig. 5A). When an update cell is fully packed with records, or as the minimum bandwidth pacing interval is reached, the update cell is transmitted to the upstream element 12.

Once received upstream, the UP 16 receives the connection indicator from the update record to identify the transmitter connection, and extracts the Fwd_Counter 38 value or the Rx_Counter 40 minus Buffer_Counter 32 value from the update record. BS_Counter 22 is reset to the value of Tx_Counter 26 minus the update record value (Fig. 7A). If this connection was disabled from transmitting due to BS_Counter 22 being equal to or greater than BS_Limit 24, this condition should now be reversed, and if so the connection should again be enabled for transmitting.

In summary, the update event provides the transmitting element 12 with an indication of how many cells originally transmitted by it have now been released from buffers within the receiving element 14, and thus provides the transmitting element 12 with a more accurate indication of receiver element 14 buffer 28 availability for that connection 20.

The buffer state check event serves two purposes: 1) it provides a mechanism to calculate and compensate for cell loss or cell insertion due to transmission errors; and 2) it

- 12 -

provides a mechanism to start (or restart) a flow if update cells were lost or if enough data cells were lost that N2_Limit 34 is never reached.

5 One timer (not shown) in the UP subsystem 16 serves all connections. The connections are enabled or disabled on a per connection basis as to whether to send check cells from the upstream transmitter element 12 to the downstream receiver element 14. The check process in the transmitter element 12 involves searching all of the connection
10 descriptors to find one which is check enabled (see Figs. 8A, 9A). Once a minimum pacing interval has elapsed (the check interval), the check cell is forwarded to the receiver element 14 and the next check enabled connection is identified. The spacing between check cells for the same
15 connection is a function of the number of active flow-controlled connections times the mandated spacing between check cells for all connections. Check cells have priority over update cells.

The check event occurs as follows, with regard to Figs.
20 8A through 8C and 9A through 9C. Each transmit element 12 connection 20 is checked after a timed check interval is reached. If the connection is flow-control enabled and the connection is valid, then a check event is scheduled for transmission to the receiver element 14. A buffer state
25 check cell is generated using the Tx_Counter 26 value for that connection 20 in the check cell payload, and is transmitted using the connection indicator from the respective connection descriptor (Figs. 8A and 9A).

In the first embodiment, a calculation of errored cells
30 is made at the receiver element 14 by summing Fwd_Counter 38 with Buffer_Counter 32, and subtracting this value from the contents of the transmitted check cell record, the value of Tx_Counter 26 (Fig. 9B). The value of Fwd_Counter 38 is increased by the errored cell count. An update record with
35 the new value for Fwd_Counter 38 is then generated. This updated Fwd_Counter 38 value subsequently updates the

- 13 -

BS_Counter 22 value in the transmitter element 12.

In the second embodiment, illustrated in Fig. 8B, the same is accomplished by resetting the Rx_Counter 40 value equal to the check cell payload value (Tx_Counter 26). A subsequent update record is established using the difference between the values of Rx_Counter 40 and Buffer_Counter 32.

Thus, the check event enables accounting for cells transmitted by the transmitter element 12, through the connection 20, but either dropped or not received by the receiver element 14.

A "no cell loss" guarantee is enabled using buffer state accounting at the connection level since the transmitter element 12 has an up-to-date account of the number of buffers 28 in the receiver element 14 available for receipt of data cells, and has an indication of when data cell transmission should be ceased due to the absence of available buffers 28 downstream.

In order to augment the foregoing protocol with a receiver element buffer sharing mechanism, link-level flow control, also known as link-level buffer state accounting, is added to connection-level flow control. It is possible for such link-level flow control to be implemented without connection-level flow control. However, a combination of the two is preferable since without connection-level flow control there would be no restriction on the number of buffers a single connection might consume.

It is desirable to perform buffer state accounting at the link level, in addition to the connection level, for the following reasons. Link-level flow control enables cell buffer sharing at a receiver element while maintaining the "no cell loss" guarantee afforded by connection-level flow control. Buffer sharing results in the most efficient use of a limited number of buffers. Rather than provide a number of buffers equal to bandwidth times RTT for each connection, a smaller number of buffers is employable in the receiver element 14 since not all connections require a full

- 14 -

compliment of buffers at any one time.

A further benefit of link-level buffer state accounting is that each connection is provided with an accurate representation of downstream buffer availability without necessitating increased reverse bandwidth for each connection. A high-frequency link-level update does not significantly effect overall per-connection bandwidth.

Link-level flow control is described now with regard to Fig. 2. Like elements found in Fig. 1 are given the same reference numbers in Fig. 2, with the addition of a prime. Once again, only one virtual connection 20' is illustrated in the link 10', though the link 10' would normally host multiple virtual connections 20'. Once again, the link 10' is a physical link in a first embodiment, and a logical grouping of plural virtual connections in a second embodiment.

The upstream transmitter element 12' (FSPP subsystem) partially includes a processor labelled From Switch Port Processor (FSPP) 16'. The FSPP processor 16' is provided with two buffer state counters, BS_Counter 22' and BS_Limit 24', and a Tx_Counter 26' each having the same function on a per-connection basis as those described with respect to Fig. 1.

The embodiment of Fig. 2 further includes a set of resources added to the upstream and downstream elements 12', 14' which enable link-level buffer accounting. These resources provide similar functions as those utilized on a per-connection basis, yet they operate on the link level.

For instance, Link_BS_Counter 50 tracks all cells in flight between the FSPP 16' and elements downstream of the receiver element 14', including cells in transit between the transmitter 12' and the receiver 14' and cells stored within receiver 14' buffers 28'. As with the update event described above with respect to connection-level buffer accounting, Link_BS_Counter 50 is modified during a link update event by subtracting either the Link_Fwd_Counter 68 value or the

- 15 -

difference between Link_Rx_Counter 70 and Link_Buffer_Counter 62 from the Link_TX_Counter 54 value. In a first embodiment, the link-level counters are implemented in external RAM associated with the FSPP processor 16'.

5 Link_BS_Limit 52 limits the number of shared downstream cell buffers 28' in the receiver element 14' to be shared among all of the flow-control enabled connections 20'. In a first embodiment, Link_BS_Counter 50 and Link_BS_Limit 52 are both twenty bits wide.

10 Link_TX_Counter 54 tracks all cells transmitted onto the link 10'. It is used during the link-level update event to calculate a new value for Link_BS_Counter 50. Link_TX_Counter 54 is twenty-eight bits wide in the first embodiment.

15 In the downstream element 14', To Switch Port Processor (TSPP) 18' also manages a set of counters for each link 10' in the same fashion with respect to the commonly illustrated counters in Figs. 1 and 2. The TSPP 18' further includes a Link_Buffer_Limit 60 which performs a function in the
20 downstream element 14' similar to Link_BS_Limit 52 in the upstream element 12' by indicating the maximum number of cell buffers 28' in the receiver 14' available for use by all connections 10'. In most cases, Link_BS_Limit 52 is equal to Link_Buffer_Limit 60. The effect of adjusting the number
25 of buffers 28' available up or down on a link-wide basis is the same as that described above with respect to adjusting the number of buffers 28 available for a particular connection 20. Link_Buffer_Limit 60 is twenty bits wide in the first embodiment.

30 Link_Buffer_Counter 62 provides an indication of the number of buffers in the downstream element 14' which are currently being used by all connections for the storage of data cells. This value is used in a check event to correct the Link_Fwd_Counter 68 (described subsequently). The
35 Link_Buffer_Counter 62 is twenty bits wide in the first embodiment.

- 16 -

Link_N2_Limit 64 and Link_N2_Counter 66, each eight bits wide in the first embodiment, are used to generate link update records, which are intermixed with connection-level update records. Link_N2_Limit 64 establishes a threshold number for triggering the generation of a link-level update record (Figs. 5B and 6B), and Link_N2_Counter 66 and Link_Fwd_Counter 68 are incremented each time a cell is released out of a buffer cell in the receiver element 14'. In a first embodiment, N2_Limit 34' and Link_N2_Limit 64 are both static once initially configured.

However, in a further embodiment of the present invention, each is dynamically adjustable based upon measured bandwidth. For instance, if forward link bandwidth is relatively high, Link_N2_Limit 64 could be adjusted down to cause more frequent link-level update record transmission. Any forward bandwidth impact would be considered minimal. Lower forward bandwidth would enable the raising of Link_N2_Limit 64 since the unknown availability of buffers 28' in the downstream element 14' is less critical.

Link_Fwd_Counter 68 tracks all cells released from buffer cells 28' in the receiver element 14' that came from the link 10' in question. It is twenty-eight bits wide in a first embodiment, and is used in the update event to recalculate Link_BS_Counter 50.

Link_Rx_Counter 70 is employed in an alternative embodiment in which Link_Fwd_Counter 68 is not employed. It is also twenty-eight bits wide in an illustrative embodiment and tracks the number of cells received across all connections 20' in the link 10'.

With regard to Figs. 2 et seq., a receiver element buffer sharing method is described. Normal data transfer by the FSPP 16' in the upstream element 12' to the TSPP 18' in the downstream element 14' is enabled across all connections 20' in the link 10' as long as the Link_BS_Counter 50 is less than or equal to Link_BS_Limit 52, as in Fig. 3B. This test prevents the FSPP 16' from transmitting more data cells than

- 17 -

it believes are available in the downstream element 14'. The accuracy of this belief is maintained through the update and check events, described next.

5 A data cell is received at the downstream element 14' if neither connection-level or link-level buffer limit are exceeded (Fig. 3B). If a limit is exceeded, the cell is discarded.

10 The update event at the link level involves the generation of a link update record when the value in Link_N2_Counter 66 reaches (equals or exceeds) the value in Link_N2_Limit 64, as shown in Figs. 5B and 6B. In a first embodiment, Link_N2_Limit 64 is set to forty.

15 The link update record, the value taken from Link_Fwd_Counter 68 in the embodiment of Fig. 6B, is mixed with the per-connection update records (the value of Fwd_Counter 38') in update cells transferred to the FSPP 16'. In the embodiment of Fig. 5B, the value of Link_Rx_Counter 70 minus Link_Buffer_Counter 62 is mixed with the per-connection update records. When the upstream element 12' receives the update cell having the link update record, it sets the Link_BS_Counter 50 equal to the value of Link_Tx_Counter 54 minus the value in the update record (Fig. 7B). Thus, Link_BS_Counter 50 in the upstream element 12' is reset to reflect the number of data cells transmitted by the upstream element 12', but not yet released in the downstream element 14'.

25 The actual implementation of the transfer of an update record, in a first embodiment, recognizes that for each TSPP subsystem 14', there is an associated FSPP processor (not illustrated), and for each FSPP subsystem 12', there is also an associated TSPP processor (not illustrated). Thus, when an update record is ready to be transmitted by the TSPP subsystem 14' back to the upstream FSPP subsystem 12', the TSPP 18' conveys the update record to the associated FSPP (not illustrated), which constructs an update cell. The cell is conveyed from the associated FSPP to the TSPP (not

- 18 -

illustrated) associated with the upstream FSPP subsystem 12'. The associated TSPP strips out the update record from the received update cell, and conveys the record to the upstream FSPP subsystem 12'.

5 The check event at the link level involves the transmission of a check cell having the Link_Tx_Counter 54 value by the FSPP 16' every "W" check cells (Figs. 8A and 9A). In a first embodiment, W is equal to four. At the receiver element 14', the TSPP 18' performs the previously described check functions at the connection-level, as well as increasing the Link_Fwd_Counter 68 value by an amount equal to the check record contents, Link_Tx_Counter 54, minus the sum of Link_Buffer_Counter 62 plus Link_Fwd_Counter 68 in the embodiment of Fig. 9C. In the embodiment of Fig. 8C, 10 Link_Rx_Counter 70 is modified to equal the contents of the check record (Link_Tx_Counter 54). This is an accounting for errored cells on a link-wide basis. An update record is then generated having a value taken from the updated Link_Fwd_Counter 68 or Link_Rx_Counter 70 values (Figs. 8C and 9C). 15 20

It is necessary to perform the check event at the link level in addition to the connection level in order to readjust the Link_Fwd_Counter 68 value (Fig. 9C) or Link_Rx_Counter 70 value (Fig. 8C) quickly in the case of large transient link failures. 25

Again with regard to Fig. 2, the following are exemplary initial values for the illustrated counters in an embodiment having 100 connections in one link.

30	BS_Limit	(24') = 20
	Buffer_Limit	(30') = 20
	N2_Limit	(34') = 3
	Link_BS_Limit	(52) = 1000
	Link_Buffer_Limit	(60) = 1000
35	Link_N2_Counter	(66) = 40

The BS_Limit value equals the Buffer_Limit value for both the connections and the link. Though BS_Limit 24' and

- 19 -

Buffer_Limit 30' are both equal to twenty, and there are 100 connections in this link, there are only 1000 buffers 28' in the downstream element, as reflected by Link_BS_Limit 52 and Link_Buffer_Limit 60. This is because of the buffer pool sharing enabled by link-level feedback.

Link-level flow control can be disabled, should the need arise, by not incrementing: Link_BS_Counter; Link_N2_Counter; and Link_Buffer_Counter, and by disabling link-level check cell transfer. No updates will occur under these conditions.

The presently described invention can be further augmented with a dynamic buffer allocation scheme, such as previously described with respect to N2_Limit 34 and Link_N2_Limit 64. This scheme includes the ability to dynamically adjust limiting parameters such as BS_Limit 24, Link_BS_Limit 52, Buffer_Limit 30, and Link_Buffer_Limit 60, in addition to N2_Limit 34 and Link_N2_Limit 64. Such adjustment is in response to measured characteristics of the individual connections or the entire link in one embodiment, and is established according to a determined priority scheme in another embodiment. Dynamic buffer allocation thus provides the ability to prioritize one or more connections or links given a limited buffer resource.

The Link_N2_Limit is set according to the desired accuracy of buffer accounting. On a link-wide basis, as the number of connections within the link increases, it may be desirable to decrease Link_N2_Limit in light of an increased number of connections in the link, since accurate buffer accounting allows greater buffer sharing among many connections. Conversely, if the number of connections within the link decreases, Link_N2_Limit may be increased, since the criticality of sharing limited resources among a relatively small number of connections is decreased.

In addition to adjusting the limits on a per-link basis, it may also be desirable to adjust limits on a per-connection basis in order to change the maximum sustained bandwidth for the connection.

- 20 -

The presently disclosed dynamic allocation schemes are implemented during link operation, based upon previously prescribed performance goals.

5 In a first embodiment of the present invention, incrementing logic for all counters is disposed within the FSPP processor 16'. Related thereto, the counters previously described as being reset to zero and counting up to a limit can be implemented in a further embodiment as starting at the limit and counting down to zero. The transmitter and receiver processors interpret the limits as starting points for the respective counters, and decrement upon detection of the appropriate event. For instance, if Buffer_Counter (or Link_Buffer_Counter) is implemented as a decrementing counter, each time a data cell is allocated to a buffer within the receiver, the counter would decrement. When a data cell is released from the respective buffer, the counter would increment. In this manner, the counter reaching zero would serve as an indication that all available buffers have been allocated. Such implementation is less easily employed in a dynamic bandwidth allocation scheme since dynamic adjustment of the limits must be accounted for in the non-zero counts.

20 A further enhancement of the foregoing zero cell loss, link-level flow control technique includes providing a plurality of shared cell buffers 28" in a downstream element 14" wherein the cell buffers 28" are divided into N prioritized cell buffer subsets, Priority 0 108a, Priority 1 108b, Priority 2 108c, and Priority 3 108d, by N - 1 threshold level(s), Threshold(1) 102, Threshold(2) 104, and Threshold(3) 106. Such a cell buffer pool 28" is illustrated in Fig. 10, in which four priorities labelled Priority 0 through Priority 3 are illustrated as being defined by three thresholds labelled Threshold(1) through Threshold(3).

30 This prioritized buffer pool enables the transmission of high priority connections while lower priority connections are "starved" or prevented from transmitting cells downstream

- 21 -

during periods of link congestion. Cell priorities are identified on a per-connection basis. The policy by which the thresholds are established is defined according to a predicted model of cell traffic in a first embodiment, or, in an alternative embodiment, is dynamically adjusted. Such dynamic adjustment may be in response to observed cell traffic at an upstream transmitting element, or according to empirical cell traffic data as observed at the prioritized buffer pool in the downstream element. For example, in an embodiment employing dynamic threshold adjustment, it may be advantageous to lower the number of buffers available to data cells having a priority less than Priority 0, or conversely to increase the number of buffers above Threshold(3), if a significantly larger quantity of Priority 0 traffic is detected.

The cell buffer pool 28" depicted in Fig. 10 is taken from the vantage point of a modified version 12" of the foregoing link-level flow control upstream element 12', the pool 28" being resident within a corresponding downstream element 14". This modified upstream element 12", viewed in Fig. 11, has at least one Link_BS_Threshold(n) 100, 102, 104 established in association with a Link_BS_Counter 50" and Link_BS_Limit 52", as described above, for characterizing a cell buffer pool 28" in a downstream element 14". These Link_BS_Thresholds 102, 104, 106 define a number of cell buffers in the pool 28" which are allocatable to cells of a given priority, wherein the priority is identified by a register 108 associated with the BS_Counter 22" counter and BS_Limit 24" register for each connection 20". The Priorities 108a, 108b, 108c, 108d illustrated in Fig. 11 are identified as Priority 0 through Priority 3, Priority 0 being the highest. When there is no congestion, as reflected by Link_BS_Counter 50" being less than Link_BS_Threshold(1) 102 in Figs. 10 and 11, flow-controlled connections of any priority can transmit. As congestion occurs, as indicated by an increasing value in the Link_BS_Counter 50", lower

- 22 -

priority connections are denied access to downstream buffers, in effect disabling their transmission of cells. In the case of severe congestion, only cells of the highest priority are allowed to transmit. For instance, with respect again to Fig. 10, only cells of Priority 0 108a are enabled for transmission from the upstream element 12" to the downstream element 14" if the link-level Link_BS_Threshold(3) 106 has been reached downstream. Thus, higher priority connections are less effected by the state of the network because they have first access to the shared downstream buffer pool. Note, however, that connection-level flow control can still prevent a high-priority connection from transmitting, if the path that connection is intended for is severely congested.

As above, Link_BS_Counter 50" is periodically updated based upon a value contained within a link-level update record transmitted from the downstream element 14" to the upstream element 12". This periodic updating is required in order to ensure accurate function of the prioritized buffer access of the present invention. In an embodiment of the present invention in which the Threshold levels 102, 104, 106 are modified dynamically, either as a result of tracking the priority associated with cells received at the upstream transmitter element or based upon observed buffer usage in the downstream receiver element, it is necessary for the FSPP 16" to have an accurate record of the state of the cell buffers 28", as afforded by the update function.

The multiple priority levels enable different categories of service, in terms of delay bounds, to be offered within a single quality of service. Within each quality of service, highest priority to shared buffers is typically given to connection/network management traffic, as identified by the cell header. Second highest priority is given to low bandwidth, small burst connections, and third highest for bursty traffic. With prioritization allocated as described, congestion within any one of the service categories will not prevent connection/management traffic from having the lowest

- 23 -

cell delay.

Initialization of the upstream element 12" as depicted in Fig. 11 is illustrated in Fig. 12A. Essentially, the same counters and registers are set as viewed in Fig. 3A for an upstream element 12" not enabling prioritized access to a shared buffer resource, with the exception that Link_BS_Threshold 102, 104, 106 values are initialized to a respective buffer value T. As discussed, these threshold buffer values can be pre-established and static, or can be adjusted dynamically based upon empirical buffer usage data.

Fig. 12B represents many of the same tests employed prior to forwarding a cell from the upstream element 12" to the downstream element 14" as shown in Fig. 3B, with the exception that an additional test is added for the provision of prioritized access to a shared buffer resource. Specifically, the FSPP 16" uses the priority value 108 associated with a cell to be transferred to determine a threshold value 102, 104, 106 above which the cell cannot be transferred to the downstream element 14". Then, a test is made to determine whether the Link_BS_Counter 50" value is greater than or equal to the appropriate threshold value 102, 104, 106. If so, the data cell is not transmitted. Otherwise, the cell is transmitted and connection-level congestion tests are executed, as previously described.

In alternative embodiments, more or less than four priorities can be implemented with the appropriate number of thresholds, wherein the fewest number of priorities is two, and the corresponding fewest number of thresholds is one. For every N priorities, there are N - 1 thresholds.

In yet a further embodiment, flow-control is provided solely at the link level, and not at the connection level, though it is still necessary for each connection to provide some form of priority indication akin to the priority field 108 illustrated in Fig. 11.

The link level flow controlled protocol as previously described can be further augmented in yet another embodiment

- 24 -

to enable a guaranteed minimum cell rate on a per-connection basis with zero cell loss. This minimum cell rate is also referred to as guaranteed bandwidth. The connection can be flow-controlled below this minimum, allocated rate, but only by the receiver elements associated with this connection. Therefore, the minimum rate of one connection is not affected by congestion within other connections.

It is a requirement of the presently disclosed mechanism that cells present at the upstream element associated with the FSPP 116 be identified by whether they are to be transmitted from the upstream element using allocated bandwidth, or whether they are to be transmitted using dynamic bandwidth. For instance, the cells may be provided in queues associated with a list labelled "preferred," indicative of cells requiring allocated bandwidth. Similarly, the cells may be provided in queues associated with a list labelled "dynamic," indicative of cells requiring dynamic bandwidth.

In a frame relay setting, the present mechanism is used to monitor and limit both dynamic and allocated bandwidth. In a setting involving purely internet traffic, only the dynamic portions of the mechanism may be of significance. In a setting involving purely CBR flow, only the allocated portions of the mechanism would be employed. Thus, the presently disclosed method and apparatus enables the maximized use of mixed scheduling connections - those requiring all allocated bandwidth to those requiring all dynamic bandwidth, and connections therebetween.

In the present mechanism, a downstream cell buffer pool 128, akin to the pool 28' of Fig. 2, is logically divided between an allocated portion 300 and a dynamic portion 301, whereby cells identified as to receive allocated bandwidth are buffered within this allocated portion 300, and cells identified as to receive dynamic bandwidth are buffered in the dynamic portion 301. Fig. 13A shows the two portions 300, 301 as distinct entities; the allocated portion is not

- 25 -

a physically distinct block of memory, but represents a number of individual cell buffers, located anywhere in the pool 128.

5 In a further embodiment, the presently disclosed mechanism for guaranteeing minimum bandwidth is applicable to a mechanism providing prioritized access to downstream buffers, as previously described in conjunction with Figs. 10 and 11. With regard to Fig. 13B, a downstream buffer pool 228 is logically divided among an allocated portion 302 and a dynamic portion 208, the latter logically subdivided by threshold levels 202, 204, 206 into prioritized cell buffer subsets 208a-d. As with Fig. 13A, the division of the buffer pool 228 is a logical, not physical, division.

10 Elements required to implement this guaranteed minimum bandwidth mechanism are illustrated in Fig. 14, where like elements from Figs. 2 and 11 are provided with like reference numbers, added to 100 or 200. Note that no new elements have been added to the downstream element; the presently described guaranteed minimum bandwidth mechanism is transparent to the downstream element.

15 New aspects of flow control are found at both the connection and link levels. With respect first to the connection level additions and modifications, D_BS_Counter 122 highlights resource consumption by tracking the number of cells scheduled using dynamic bandwidth transmitted downstream to the receiver 114. This counter has essentially the same function as BS_Counter 22' found in Fig. 2, where there was no differentiation between allocated and dynamically scheduled cell traffic. Similarly, D_BS_Limit 124, used to provide a ceiling on the number of downstream buffers available to store cells from the transmitter 112, finds a corresponding function in BS_Limit 24' of Fig. 2. As discussed previously with respect to link level flow control, the dynamic bandwidth can be statistically shared; the actual number of buffers available for dynamic cell traffic can be over-allocated. The amount of "D" buffers

20
25
30
35

- 26 -

provided to a connection is equal to the RTT times the dynamic bandwidth plus $N2$. RTT includes delays incurred in processing the update cell.

5 A_BS_Counter 222 and A_BS_Limit 224 also track and limit, respectively, the number of cells a connection can transmit by comparing a transmitted number with a limit on buffers available. However, these values apply strictly to allocated cells; allocated cells are those identified as requiring allocated bandwidth (the guaranteed minimum bandwidth) for transmission. Limit information is set up at
10 connection initialization time and can be raised and lowered as the guaranteed minimum bandwidth is changed. If a connection does not have an allocated component, the A_BS_Limit 224 will be zero. The A_BS_Counter 222 and
15 A_BS_Limit 224 are in addition to the D_BS_Counter 122 and D_BS_Limit 124 described above. The amount of "A" buffers dedicated to a connection is equal to the RTT times the allocated bandwidth plus $N2$. The actual number of buffers dedicated to allocated traffic cannot be over-allocated.
20 This ensures that congestion on other connections does not impact the guaranteed minimum bandwidth.

 A connection loses, or runs out of, its allocated bandwidth through the associated upstream switch once it has enqueued a cell but has no more "A" buffers as reflected by
25 A_BS_Counter 222 and A_BS_Limit 224. If a connection is flow controlled below its allocated rate, it loses a portion of its allocated bandwidth in the switch until the congestion condition is alleviated. Such may be the case in multipoint-to-point (M2P) switching, where plural sources on the same
30 connection, all having a minimum guaranteed rate, converge on a single egress point which is less than the sum of the source rates. In an embodiment of the presently disclosed mechanism in which the transmitter element is a portion of a switch having complimentary switch flow control, the
35 condition of not having further "A" buffer states inhibits the intra-switch transmission of further allocated cell

- 27 -

traffic for that connection.

The per-connection buffer return policy is to return buffers to the allocated pool first, until the A_BS_Counter 222 equals zero. Then buffers are returned to the dynamic pool, decreasing D_BS_Counter 122.

Tx_Counter 126 and Priority 208 are provided as described above with respect to connection-level flow control and prioritized access.

On the link level, the following elements are added to enable guaranteed minimum cell rate on a per-connection basis. Link_A_BS_Counter 250 is added to the FSPP 116. It tracks all cells identified as requiring allocated bandwidth that are "in-flight" between the FSPP 116 and the downstream switch fabric, including cells in the TSPP 118 cell buffers 128, 228. The counter 250 is decreased by the same amount as the A_BS_Counter 222 for each connection when a connection level update function occurs (discussed subsequently).

Link_BS_Limit 152 reflects the total number of buffers available to dynamic cells only, and does not include allocated buffers. Link_BS_Counter 150, however, reflects a total number of allocated and dynamic cells transmitted. Thus, connections are not able to use their dynamic bandwidth when Link_BS_Counter 150 (all cells in-flight, buffered, or in downstream switch fabric) minus Link_A_BS_Counter 250 (all allocated cells transmitted) is greater than Link_BS_Limit 152 (the maximum number of dynamic buffers available). This is necessary to ensure that congestion does not impact the allocated bandwidth. The sum of all individual A_BS_Limit 224 values, or the total per-connection allocated cell buffer space 300, 302, is in one embodiment less than the actually dedicated allocated cell buffer space in order to account for the potential effect of stale (i.e., low frequency) connection-level updates.

Update and check events are also implemented in the presently disclosed allocated/dynamic flow control mechanism. The downstream element 114 transmits connection level update

- 28 -

cells when either a preferred list and a VBR-priority 0 list are empty and an update queue is fully packed, or when a "max_update_interval" (not illustrated) has been reached.

At the upstream end 112, the update cell is analyzed to identify the appropriate queue, the FSPP 116 adjusts the A_BS_Counter 222 and D_BS_Counter 122 for that queue, returning cell buffers to "A" first then "D", as described above, since the FSPP 116 cannot distinguish between allocated and dynamic buffers. The number of "A" buffers returned to individual connections is subtracted from Link_A_BS_Counter 250.

Other link level elements used in association with the presently disclosed minimum guaranteed bandwidth mechanism, such as Link_Tx_Counter 154, function as described in the foregoing discussion of link level flow control. Also, as previously noted, a further embodiment of the presently described mechanism functions with a link level flow control scenario incorporating prioritized access to the downstream buffer resource 228 through the use of thresholds 202, 204, 206. The function of these elements are as described in the foregoing.

The following is an example of a typical initialization in a flow controlled link according to the present disclosure:

Downstream element has 3000 buffers;
Link is short haul, so $RTT \times \text{bandwidth}$ equals one cell;
100 allocated connections requiring 7 "A"
buffers each, consuming 700 buffers
total;

$3000 - 700 = 2300$ "D" buffers to be shared among 512
connections having zero allocated bandwidth;
Link_BS_Limit = 2300.

If D_BS_Counter \geq D_BS_Limit, then the queue is prevented from indicating that it has a cell ready to transmit. In the embodiment referred to above in which the upstream element is a switch having composite bandwidth, this

- 29 -

occurs by the queue being removed from the dynamic list, preventing the queue from being scheduled for transmit using dynamic bandwidth.

For allocated cells, a check is made when each cell is
5 enqueued to determine whether the cell, plus other enqueued
 cells, plus A_BS_Counter, is a number greater than
 A_BS_Limit. If not, the cell is enqueued and the queue is
 placed on the preferred list. Else, the connection is
 prevented from transmitting further cells through the
10 upstream element 112 switch fabric.

Initialization of the upstream element 112 as depicted
in Fig. 14 is illustrated in Fig. 15A. Essentially, the same
counters and registers set in Fig. 3A for an upstream element
12' (when prioritized access to a shared buffer resource is
15 not enabled), and in Fig. 12A for an upstream element 12"
 (when prioritized access is enabled). Exceptions include:
 Link_A_BS_Counter 250 initialized to zero; connection-level
 allocated and dynamic BS_Counters 122, 222 set to zero; and
 connection-level allocated and dynamic BS_Limits 124, 224 set
20 to respective values of N_A and N_B . Similarly, on the
 downstream end at the connection level, the allocated and
 dynamic Buffer_Limits and Buffer_counters are set, with the
 Buffer_Limits employing a bandwidth value for the respective
 traffic type (i.e., BW_A = allocated cell bandwidth and BW_D =
25 dynamic cell bandwidth). Further, each cell to be
 transmitted is identified as either requiring allocated or
 dynamic bandwidth as the cell is received from the switch
 fabric.

Fig. 15B represents many of the same tests employed
30 prior to forwarding a cell from the upstream element 112 to
 the downstream element 114 as shown in Figs. 3B and 12B, with
 the following exceptions. Over-allocation of buffer states
 per connection is checked for dynamic traffic only and is
 calculated by subtracting Link_A_BS_Counter from
35 Link_BS_Counter and comparing the result to Link_BS_Limit.
 Over-allocation on a link-wide basis is calculated from a

- 30 -

summation of Link_BS_Counter (which tracks both allocated and dynamic cell traffic) and Link_A_BS_Counter against the Link_BS_Limit. Similarly, over-allocation at the downstream element is tested for both allocated and dynamic traffic at the connection level. As previously indicated, the presently disclosed mechanism for providing guaranteed minimum bandwidth can be utilized with or without the prioritized access mechanism, though aspects of the latter are illustrated in Fig. 15A and 15B for completeness.

As discussed, connection-level flow control as known in the art relies upon discrete control of each individual connection. In particular, between network elements such as a transmitting element and a receiving element, the control is from transmitter queue to receiver queue. Thus, even in the situation illustrated in Fig. 16 in which a single queue Q_A in a transmitter element is the source of data cells for four queues Q_W , Q_X , Q_Y , and Q_Z associated with a single receiver processor, the prior art does not define any mechanism to handle this situation.

In Fig. 16, the transmitter element 10 is an FSPP element having a FSPP 11 associated therewith, and the receiver element 12 is a TSPP element having a TSPP 13 associated therewith. The FSPP 11 and TSPP 13 as employed in Fig. 16 selectively provide the same programmable capabilities as described above, such as link-level flow control, prioritized access to a shared, downstream buffer resource, and guaranteed minimum cell rate on a connection level, in addition to a connection-level flow control mechanism. Whether one or more of these enhanced capabilities are employed in conjunction with the connection-level flow control is at the option of the system configurator.

Yet another capability provided by the FSPP and TSPP according to the present disclosure is the ability to treat a group of receiver queues jointly for purposes of connection-level flow control. In Fig. 16, instead of

- 31 -

utilizing four parallel connections, the presently disclosed mechanism utilizes one connection 16 in a link 14, terminating in four separate queues Q_w , Q_x , Q_y , and Q_z , though the four queues are treated essentially as a single, joint entity for purposes of connection-level flow control. This is needed because some network elements need to use a flow controlled service but cannot handle the bandwidth of processing update cells when N2 is set to a low value, 10 or less (see above for a discussion of the update event in connection-level flow control). Setting N2 to a large value, such as 30, for a large number of connections requires large amounts of downstream buffering because of buffer orphaning, where buffers are not in-use but are accounted for up-stream as in-use because of the lower frequency of update events. This mechanism is also useful to terminate Virtual Channel Connections (VCC) within a Virtual Path Connection (VPC), where flow control is applied to the VPC.

This ability to group receiver queues is a result of manipulations of the queue descriptor associated with each of the receiver queues Q_w , Q_x , Q_y , and Q_z . With reference to Fig. 17, queue descriptors for the queues in the receiver are illustrated. Specifically, the descriptors for queues Q_w , Q_x , and Q_y are provided on the left, and in general have the same characteristics. One of the first fields pertinent to the present disclosure is a bit labelled "J." When set, this bit indicates that the associated queue is being treated as part of a joint connection in a receiver. Instead of maintaining all connection-level flow control information in each queue descriptor for each queue in the group, certain flow control elements are maintained only in one of the queue descriptors for the group. In the illustrated case, that one queue is queue Q_z .

In each of the descriptors for queues Q_w , Q_x , and Q_y , a "Joint Number" field provides an offset or pointer to a set of flow control elements in the descriptor for queue Q_z . This pointer field may provide another function when the "J"

- 32 -

bit is not set. While Buffer_Limit (labelled "Buff_Limit" in Fig. 17) and N2_Limit are maintained locally within each respective descriptor, Joint_Buffer_Counter (labelled "Jt_Buff_Cntr"), Joint_N2_Counter (labelled "Jt_N2_Cntr"), and Joint_Forward_Counter (labelled "Jt_Fwd_Cntr") are maintained in the descriptor for queue Q_z for all of the queues in the group. The same counters in the descriptors for queues Q_w , Q_x , and Q_y go unused. The joint counters perform the same function as the individual counters, such as those illustrated in Fig. 2 at the connection level, but are advanced or decremented as appropriate by actions taken in association with the individual queues. Thus, for example, Joint_Buffer_Counter is updated whenever a buffer cell receives a data cell or releases a data cell in association with any of the group queues. The same applies to Joint_N2_Counter and Joint_Forward_Counter. In an alternate embodiment of the previously described flow control mechanism, each Forward_Counter is replaced with Receive_Counter. Similarly, in an alternative embodiment of the presently disclosed mechanism, Joint_Forward_Counter is replaced with Joint_Receive_Counter, depending upon which is maintained in each of the group queues. Only the embodiment including Forward_Counter and Joint_Forward_Counter are illustrated.

Not all of the per-queue descriptor elements are superseded by functions in a common descriptor. Buffer_Limit (labelled "Buff_Limit" in Fig. 17) is set and referred to on a per-queue basis. Thus, Joint_Buffer_Counter is compared against the Buffer_Limit of a respective queue. Optionally, the Buffer_Limit could be Joint_Buffer_Limit, instead of maintaining individual, common limits. The policy is to set the same Buffer_Limit in all the TSPP queues associated with a single Joint_Buffer_Counter.

An update event is triggered, as previously described, when the Joint_N2_Counter reaches the queue-level N2_Limit. The policy is to set all of the N2_Limits equal to the same

- 33 -

value for all the queues associated with a single joint flow control connection.

When a check cell is received for a connection, an effort to modify the Receive_Counter associated with the receiving queue results in a modification of the Joint_Receive_Counter. Thus, the level of indirection provided by the Joint_Number is applicable to both data cells and check cells.

At the transmitter element 10, only one set of upstream flow control elements are maintained. At connection set-up time, the joint connection is set-up as a single, point-to-point connection, as far as the upstream elements are concerned. Therefore, instead of maintaining four sets of upstream elements for the embodiment of Fig. 16, the presently disclosed mechanism only requires one set of elements (Tx_Counter, BS_Counter, BS_Limit, all having the functionality as previously described).

Once a joint flow control entity has been established, other TSPP queues for additional connections may be added. To do so, each new queue must have the same N2_Limit and Buffer_Limit values. The queues for the additional connections will reference the common Joint_N2_Counter and either Joint_Forward_Counter or Joint_Receive_Counter.

As previously noted, when $J = 1$, the Joint_Number field is used as an offset to the group descriptor. The Joint_Number for the group descriptor is set to itself, as shown in Fig. 17 with regard to the descriptor for queue Q_z . This is also the case in point-to-point connections (VCC to VCC rather than the VPC to VCC, as illustrated in Fig. 16), where each Joint_Number points to its own descriptor. Implementation for each of point-to-point and the presently described point-to-multipoint connections is thus simplified.

Having described preferred embodiments of the invention, it will be apparent to those skilled in the art that other embodiments incorporating the concepts may be used. These and other examples of the invention illustrated above are

- 34 -

intended by way of example and the actual scope of the invention is to be determined from the following claims.

- 35 -

CLAIMS

What is claimed is:

1. A method for providing prioritized access to a shared buffer resource in a receiver apparatus by data cells transmitted over a link to a receiver apparatus by a transmitter apparatus, said method comprising the steps of:

determining, by said transmitter apparatus, a priority level of a data cell to be transmitted;

generating, in said transmitter apparatus, a first count indicative of a number of data cells transmitted to said receiver apparatus over said link for storage in said shared buffer resource;

storing, in said transmitter apparatus, at least one buffer threshold providing a maximum number of buffers allocatable to data cells of a respective priority level;

identifying, by said transmitter apparatus, which, if any, of said at least one buffer thresholds are equaled or exceeded by said first count; and

disabling transmission of said data cell to be transmitted, by said transmitter apparatus, if said at least one buffer threshold corresponding to said priority level of said data cell to be transmitted is equaled or exceeded by said first count.

2. The method according to claim 1, wherein said step of determining a priority level further comprises determining said priority level of said data cell to be transmitted based upon a connection of said data cell.

3. The method according to claim 1, wherein N - 1 buffer thresholds are stored for every N priority levels determined.

4. The method according to claim 1, further comprising the step of modifying said first count as data cells are transmitted by said transmitter apparatus to said receiver apparatus over said link.

- 36 -

5 5. The method according to claim 1, wherein said step of modifying said first count further comprises receiving, from said receiver apparatus and in said transmitter apparatus, a value identifying the number of buffers allocated to data cells within said shared buffer resource and modifying said first count in accordance therewith.

10 6. The method according to claim 1, wherein said step of storing at least one buffer threshold further comprises dynamically adjusting said at least one buffer threshold according to an empirical analysis of data cells to be transmitted by said transmitter apparatus.

15 7. The method according to claim 6, wherein said step of dynamically adjusting said at least one buffer threshold further comprises dynamically adjusting said at least one buffer threshold according to an empirical analysis, at said transmitter apparatus, of priorities of data cells to be transmitted by said transmitter apparatus.

20 8. The method according to claim 1, wherein said step of storing at least one buffer threshold further comprises dynamically adjusting said at least one buffer threshold according to an empirical analysis of utilization of said shared buffer resource by said transmitter apparatus.

25 9. The method according to claim 1, further comprising the step of establishing, in said transmitter apparatus, a maximum number of data cells storable in said shared buffer resource within said receiver apparatus.

30 10. The method according to claim 9, wherein said step of disabling transmission further comprises disabling transmission of said data cell if said first count equals or exceeds said maximum number of data cells storable.

- 37 -

11. The method according to claim 1, wherein said data cells are ATM cells.

12. The method according to claim 1, wherein said data cells are of a predetermined, fixed length.

13. A method of providing prioritized access by data cells of a transmitter element to a shared buffer resource in a receiver element via a communications resource, said transmitter element having a memory resource associated therewith, said method comprising the steps of:

defining, by said transmitter element, a maximum number of priority allocated buffers available in said shared buffer resource for each of a plurality of data cell priorities and storing each of said maximum numbers in said memory resource;

storing, in said memory resource, a value indicative of a number of data cells transmitted by said transmitter element to said receiver element via said communication resource for storage in said shared buffer resource;

modifying said value as data cells are transmitted to said shared buffer resource;

identifying, in said transmitter element, a priority assigned to a first data cell; and

transmitting said first data cell by said transmitter element to said receiver element via said communications resource for storage in said shared buffer resource if said value is less than a respective maximum number of priority allocated buffers.

14. The method according to claim 13, further comprising the step of dynamically adjusting said maximum numbers of priority allocated buffers according to an empirical analysis of data cells to be transmitted by said transmitter element.

15. The method according to claim 14, wherein said step of dynamically adjusting said maximum numbers of priority

- 38 -

allocated buffers further comprises dynamically adjusting said maximum numbers of priority allocated buffers according to an empirical analysis, at said transmitter element, of priorities of said data cells of said transmitter element.

5

16. The method according to claim 13, further comprising the step of dynamically adjusting said maximum numbers of priority allocated buffers according to an empirical analysis, at said transmitter element, of utilization of said shared buffer resource.

10

17. The method according to claim 13, wherein said step of identifying a priority further comprises identifying said priority assigned to said first data cell based upon a connection of said first data cell.

15

18. The method according to claim 13, wherein for N data cell priorities, said step of defining a maximum number further comprises:

20 defining one maximum number of priority allocated data cells equivalent to a buffer resource capacity; and
 defining N - 1 maximum number of priority allocated data cells equivalent to data cell priority thresholds.

25 19. The method according to claim 13, wherein said step of modifying said value further comprises receiving, from said receiver element and in said transmitter element, a value identifying the number of buffers allocated to data cells within said shared buffer resource and modifying said value
30 in accordance therewith.

20. The method according to claim 13, wherein said data cells are ATM cells.

35 21. The method according to claim 13, wherein said data cells are of a predetermined, fixed length.

- 39 -

22. A prioritized buffer sharing apparatus comprising:

a communications medium having a transmitter end and a receiver end;

5 a transmitter at said transmitter end of said communications medium for transmitting a data cell over said communications medium; and

10 a receiver at said receiver end of said communications medium for receiving said data cell, said receiver having a prioritized shared buffer resource associated therewith for storing said data cell,

wherein said transmitter analyzes said data cell to determine a priority associated therewith, and

15 wherein said transmitter transmits said data cell to said receiver via said communications medium for storage in said shared buffer resource if buffers in said shared buffer resource are available for storing data cells of like priority.

20 23. The prioritized buffer sharing apparatus according to claim 22, wherein said transmitter further comprises a processing element and a memory element in which is maintained, by said processing element:

25 a first value representative of a maximum number of data cells storable in said prioritized shared buffer resource;

a second value representative of a number of buffers transmitted to said prioritized shared buffer resource for storage therein; and

30 at least one threshold value each representative of a maximum number of data cells of a respective priority which are storable in said prioritized shared buffer resource.

35 24. The prioritized buffer sharing apparatus according to claim 23, wherein said processing element prohibits transmission of said data cell from said transmitter to said receiver if said second value is not less than said first

- 40 -

value.

25. The prioritized buffer sharing apparatus according to claim 23, wherein said processing element prohibits transmission of said data cell if said second value is not less than a respective threshold value for the priority of said data cell.

26. The prioritized buffer sharing apparatus according to claim 23, wherein said at least one threshold value is dynamically adjustable by said processing element based upon an empirical analysis of data cells to be transmitted to said shared buffer resource.

27. The prioritized buffer sharing apparatus according to claim 26, wherein said at least one threshold value is dynamically adjustable by said processing element based upon an empirical analysis, at said transmitter, of priorities of data cells to be transmitted to said shared buffer resource.

28. The prioritized buffer sharing apparatus according to claim 23, wherein said at least one threshold value is dynamically adjustable by said processing element based upon an empirical analysis of utilization of said shared buffer resource.

29. The prioritized buffer sharing apparatus according to claim 23, wherein there are N possible data cell priorities and N - 1 threshold values.

30. The prioritized buffer sharing apparatus according to claim 23, wherein said second value is updated by said processing element according to data received in said transmitter from said receiver reflective of a number of data cells presently stored within said shared buffer resource.

- 41 -

31. The prioritized buffer sharing apparatus according to claim 22, wherein said data cell is an ATM cell.

5 32. The prioritized buffer sharing apparatus according to claim 22, wherein said data cell is of a predetermined, fixed length.

1/23

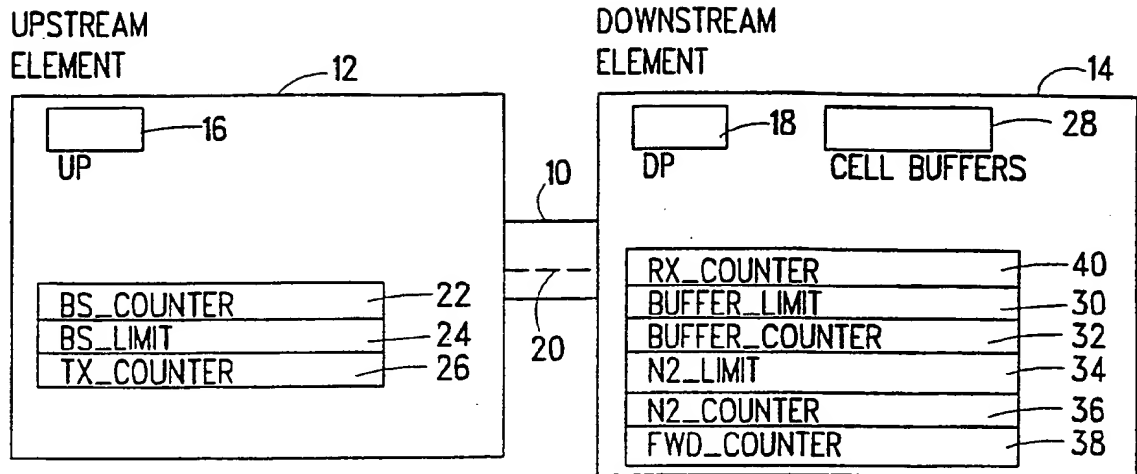


FIG. 1
PRIOR ART

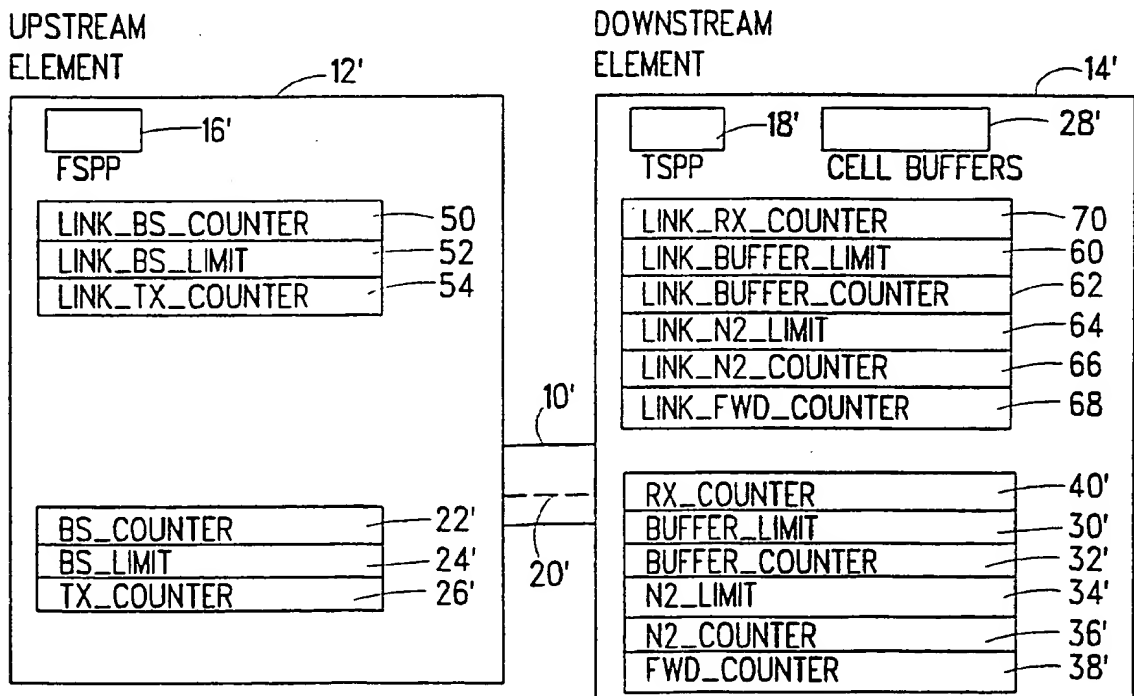


FIG. 2

2/23

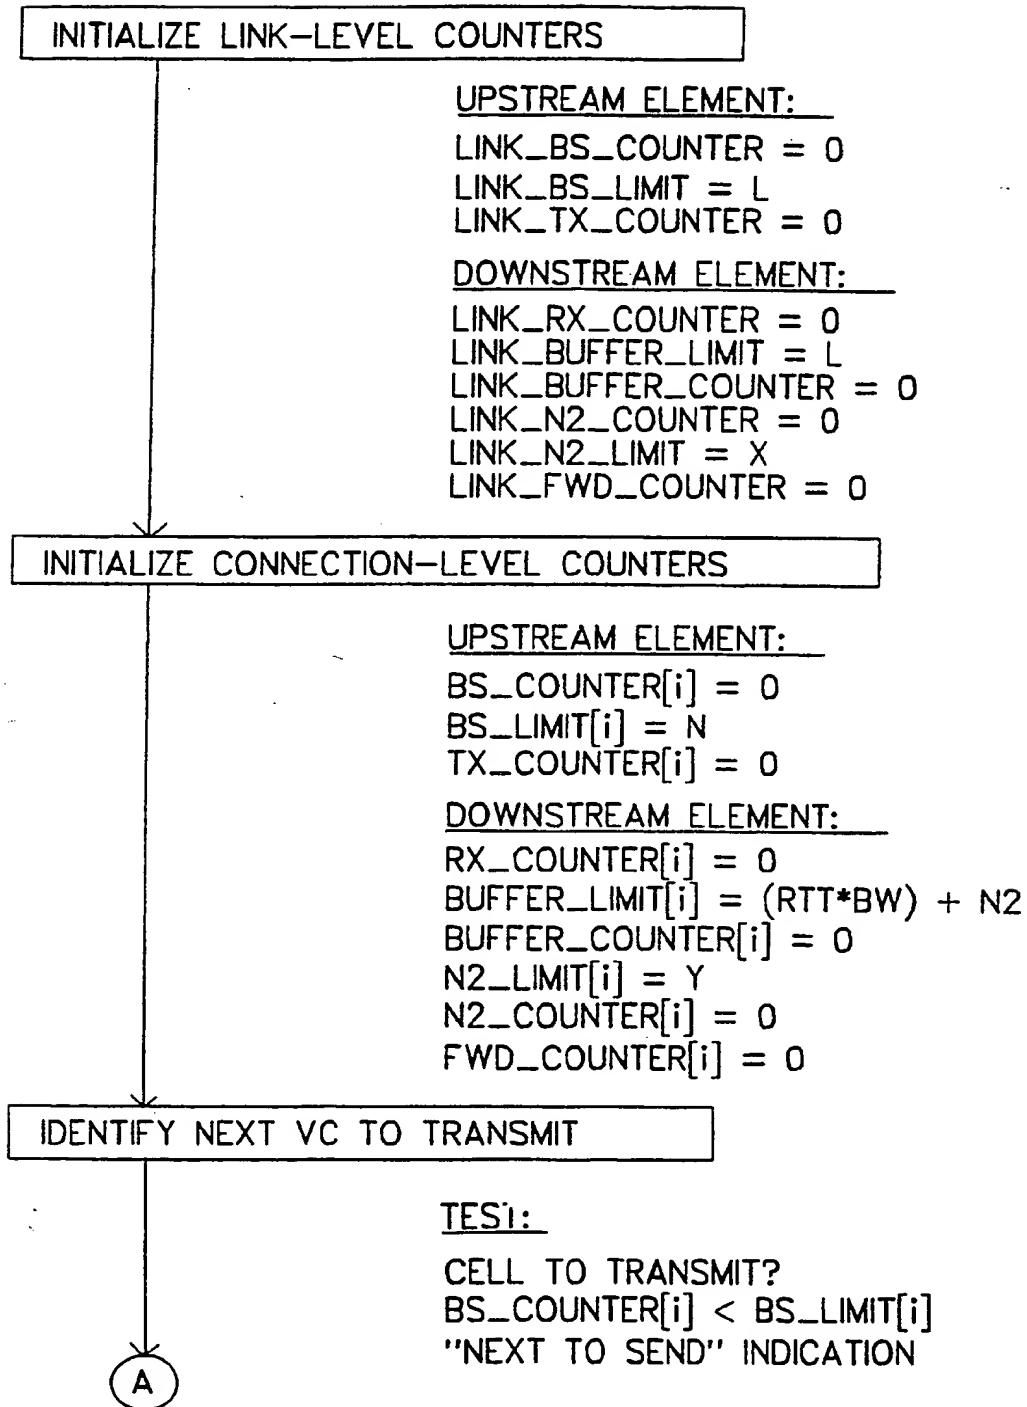


FIG. 3A

3/23

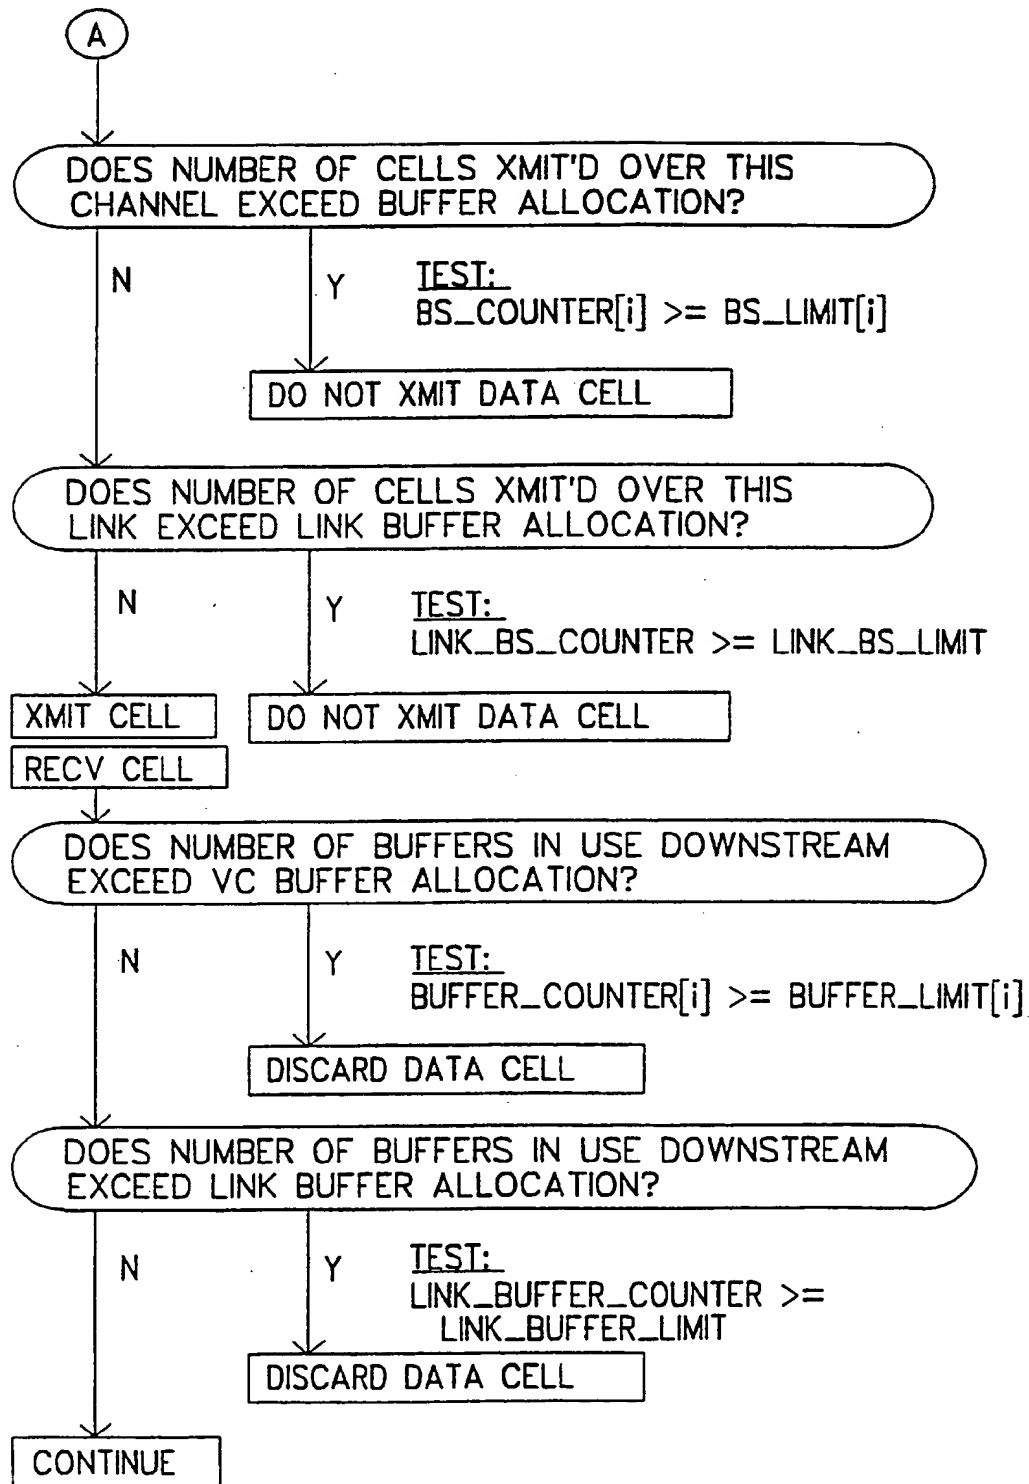


FIG. 3B

4/23

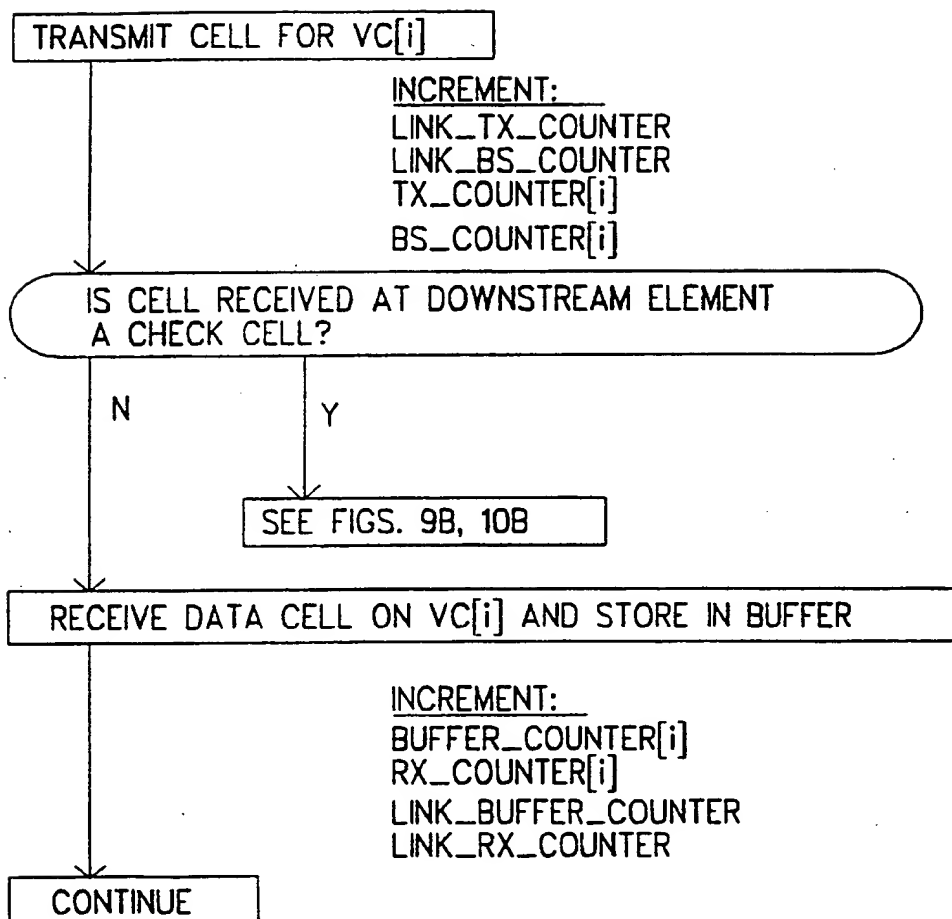


FIG. 4

5/23

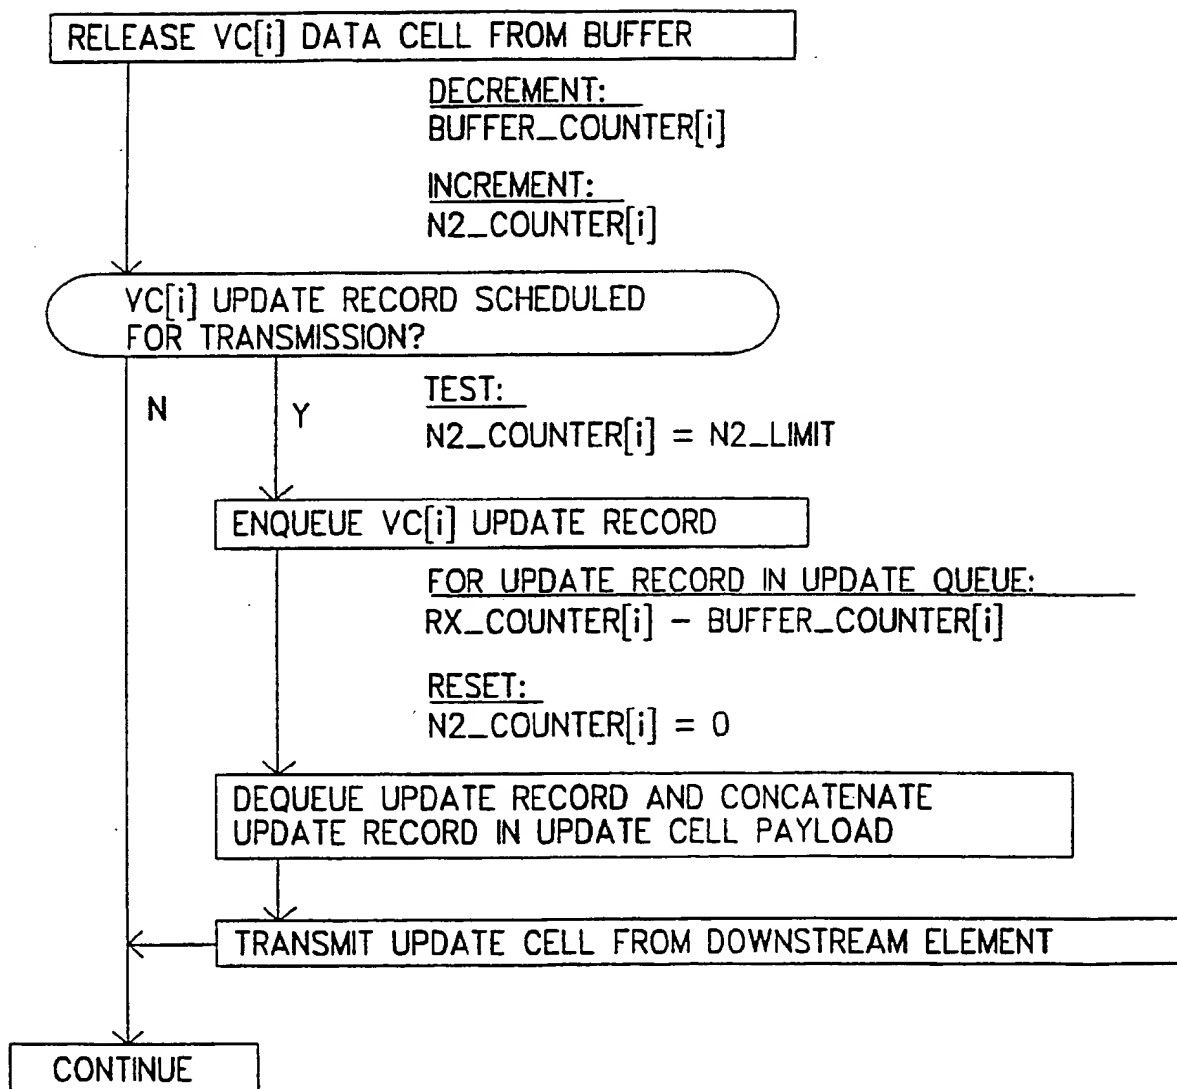


FIG. 5A

6/23

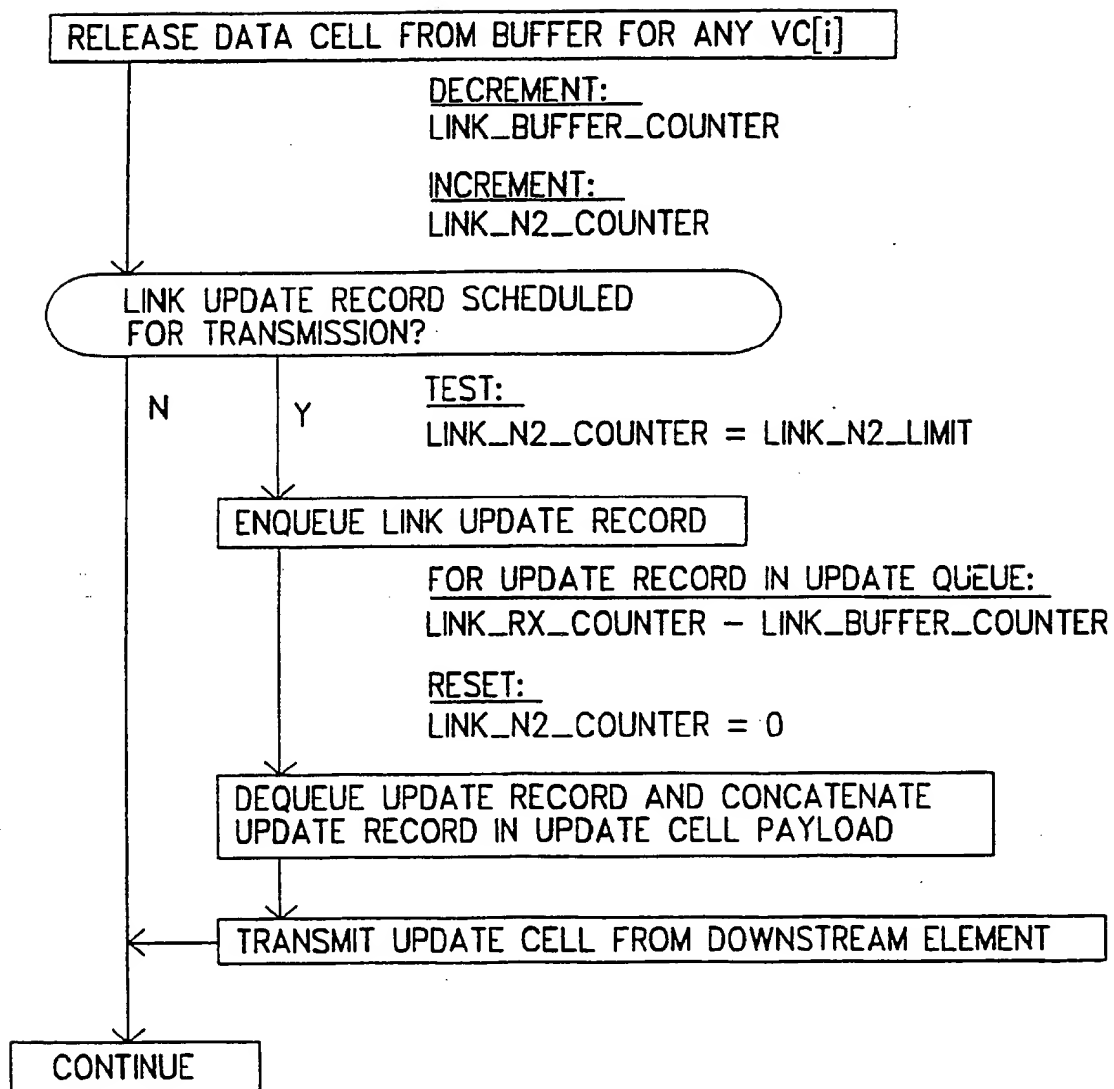


FIG. 5B

7/23

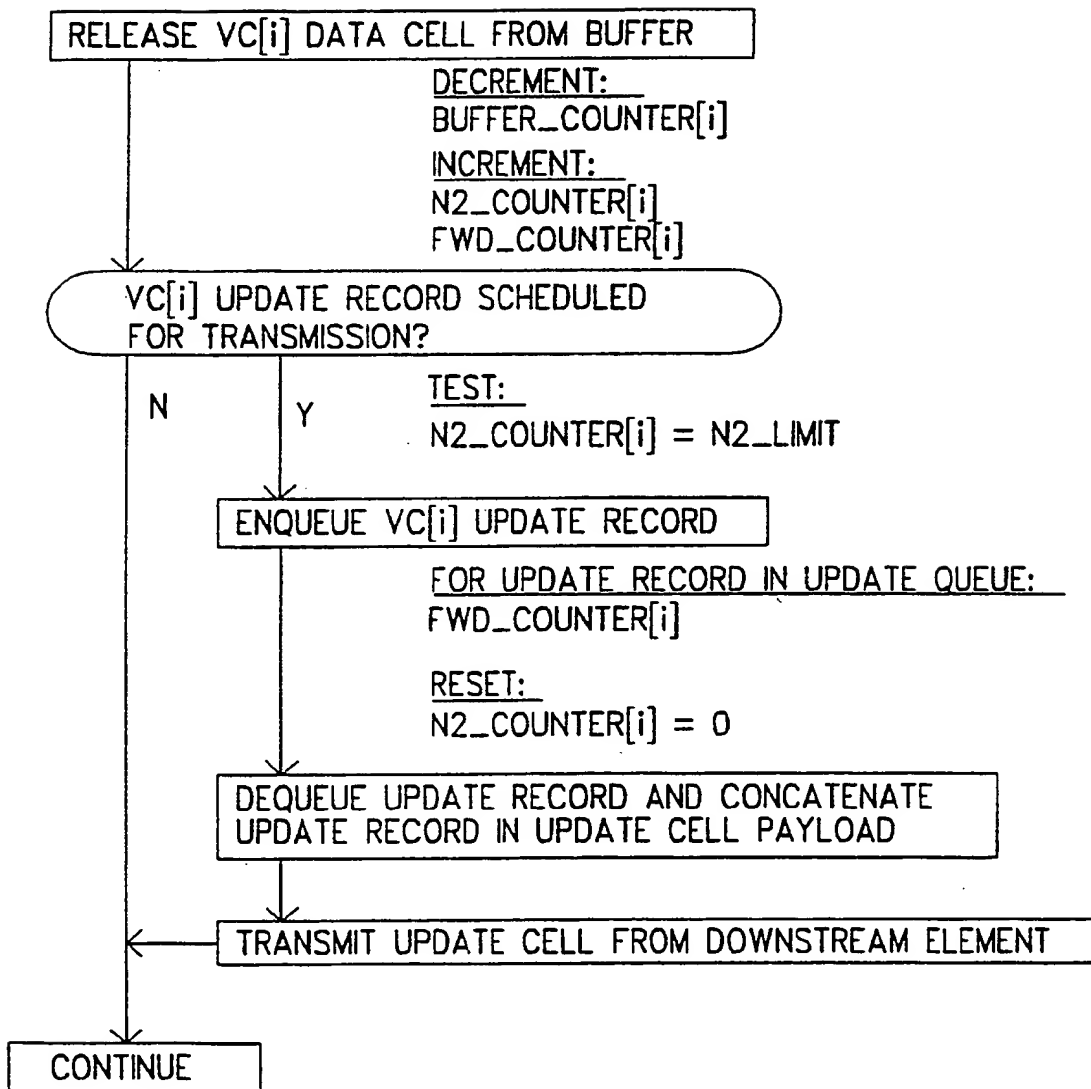


FIG. 6A

8/23

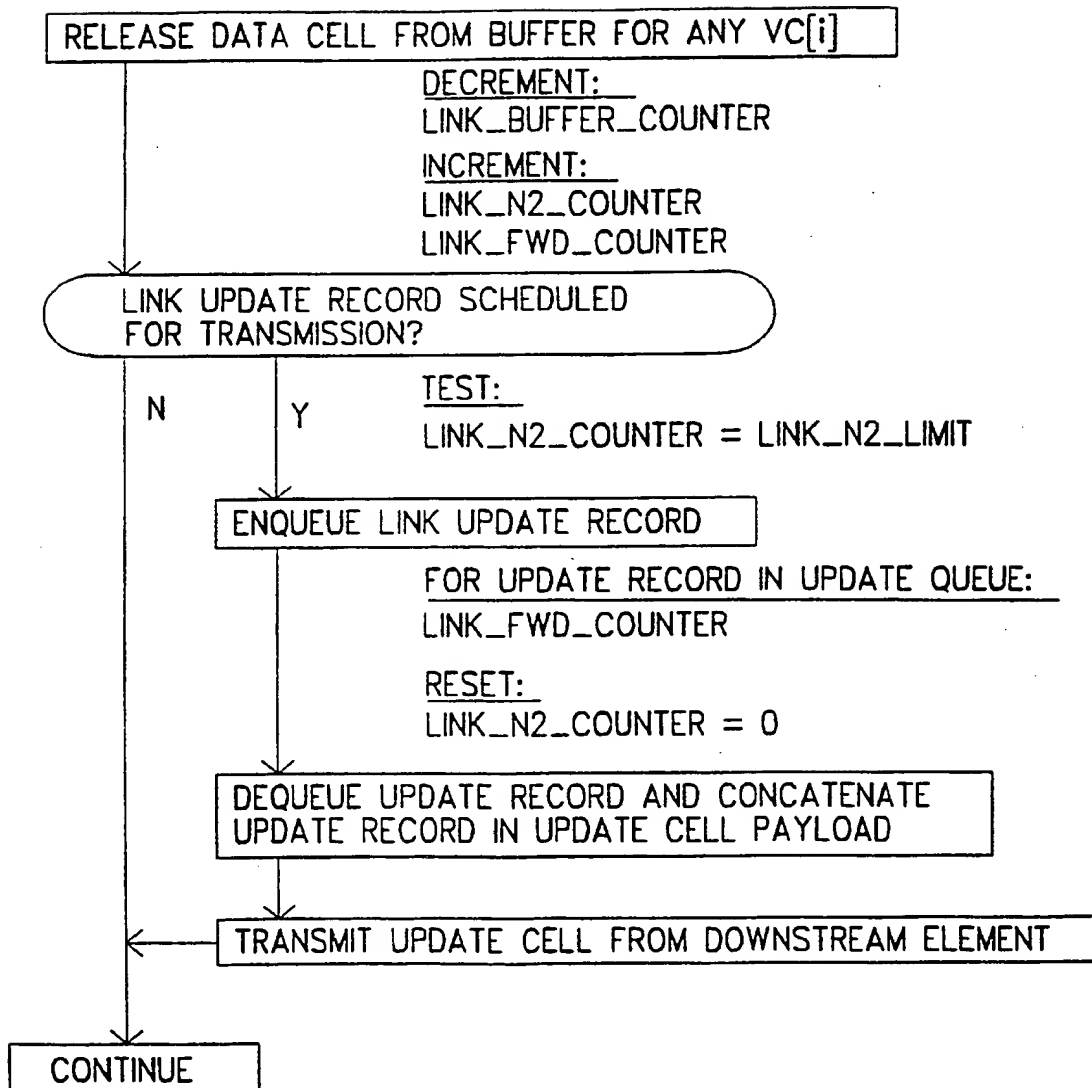


FIG. 6B

9/23

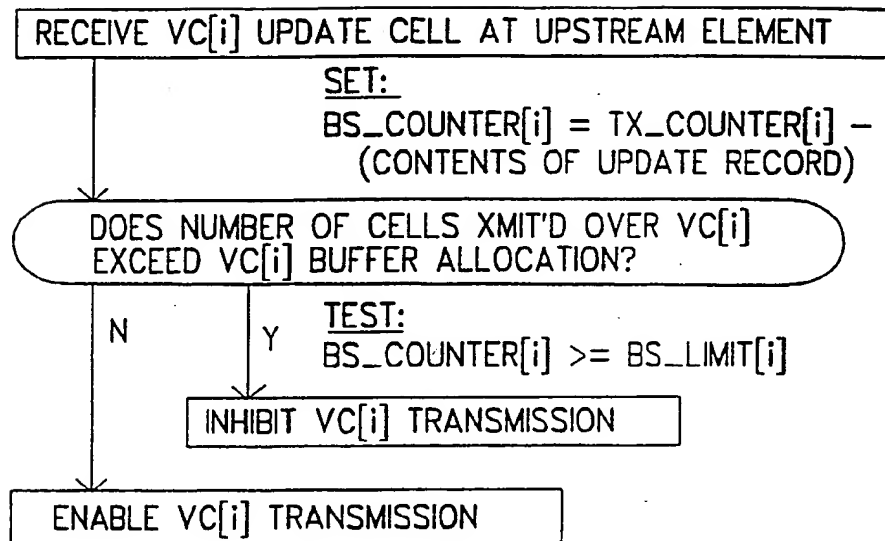


FIG. 7A

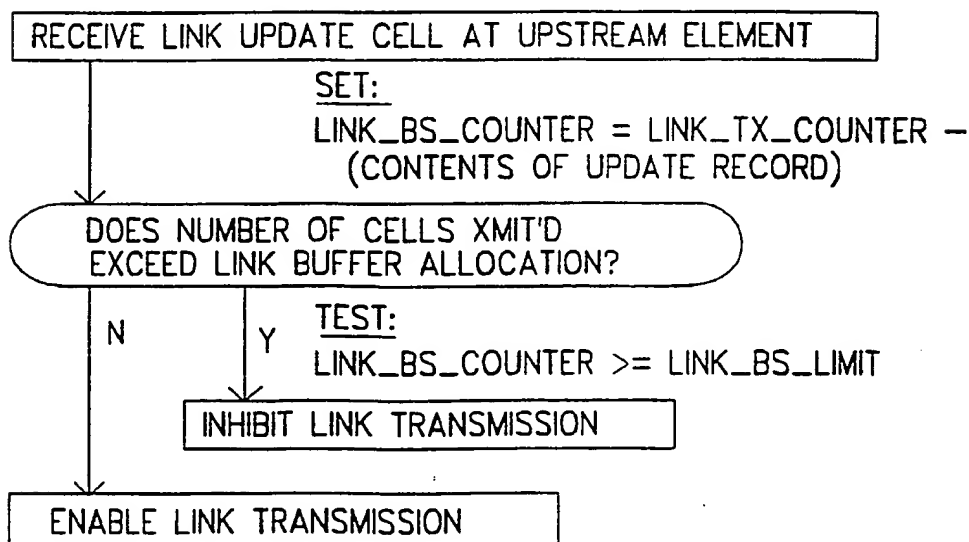


FIG. 7B

10/23

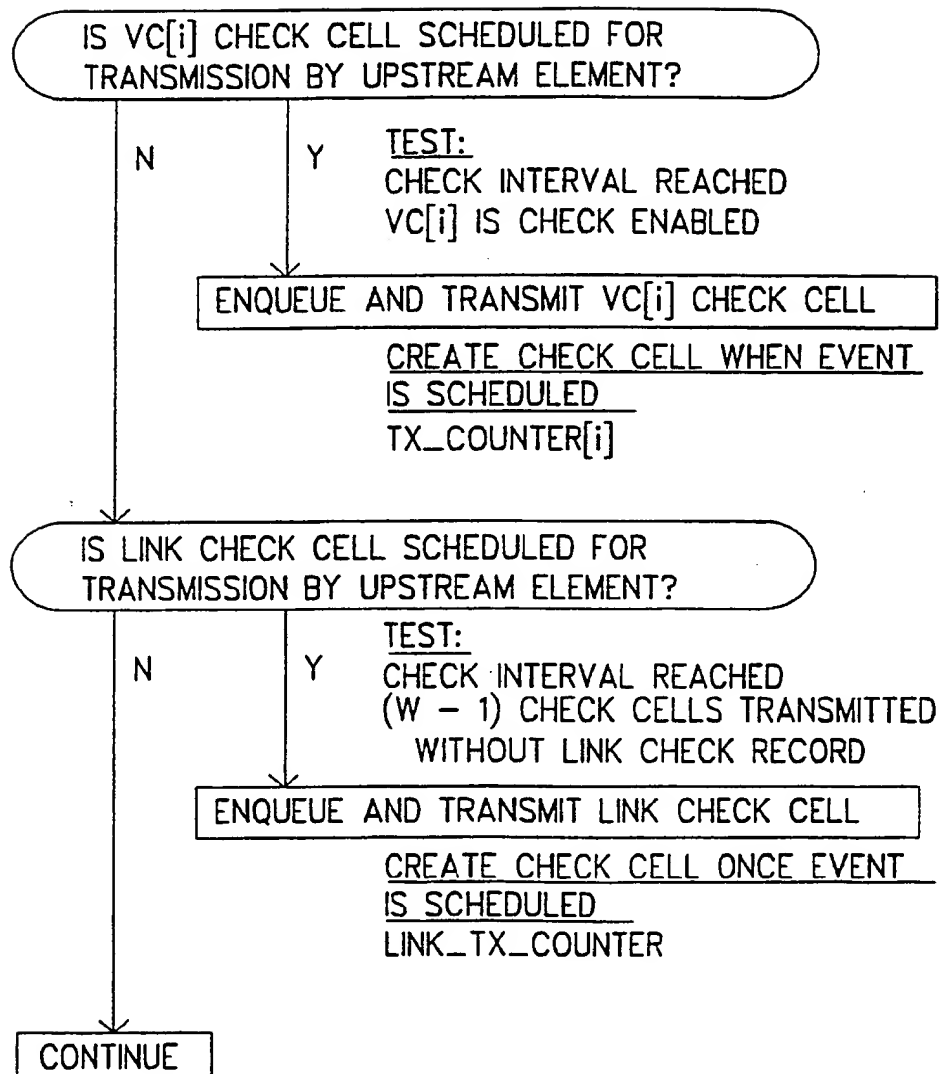


FIG. 8A

11/23

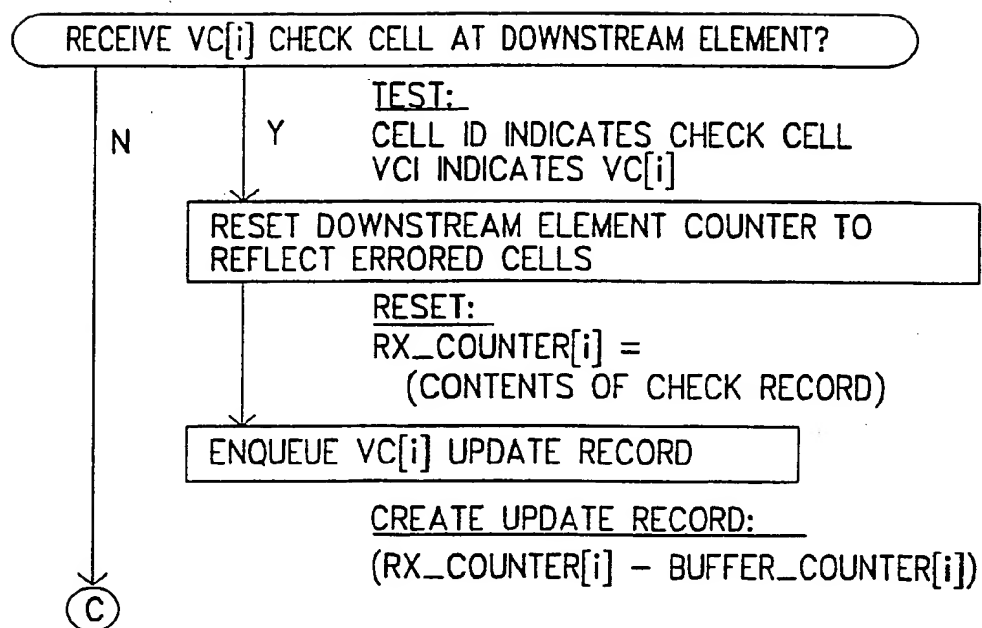


FIG. 8B

12/23

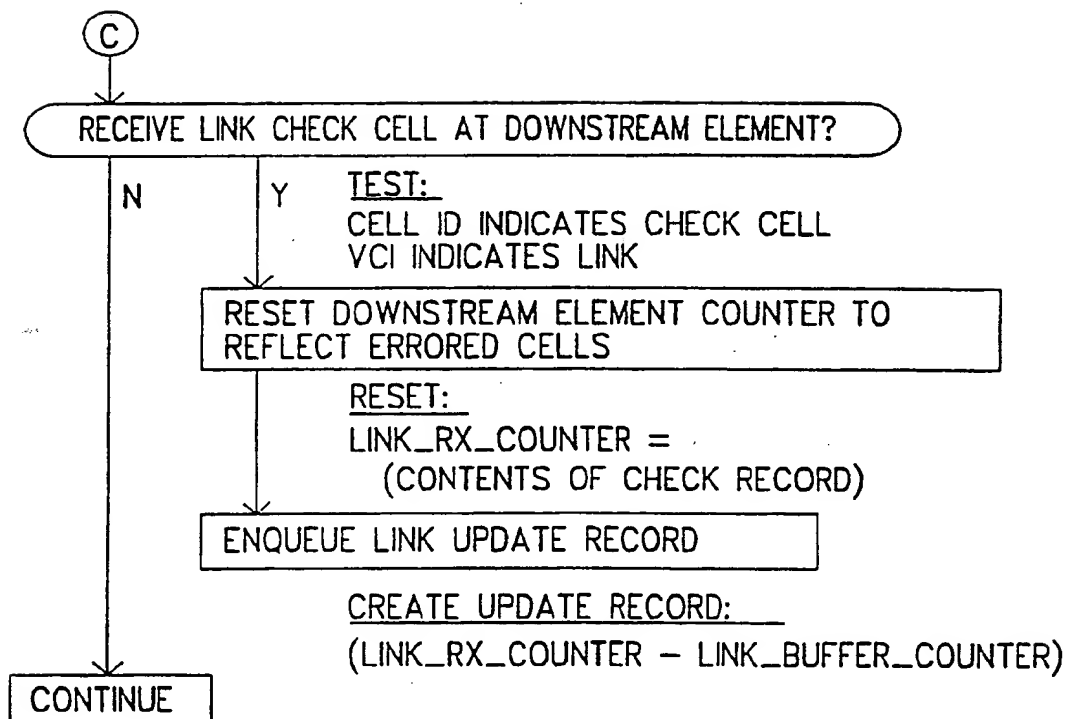


FIG. 8C

13/23

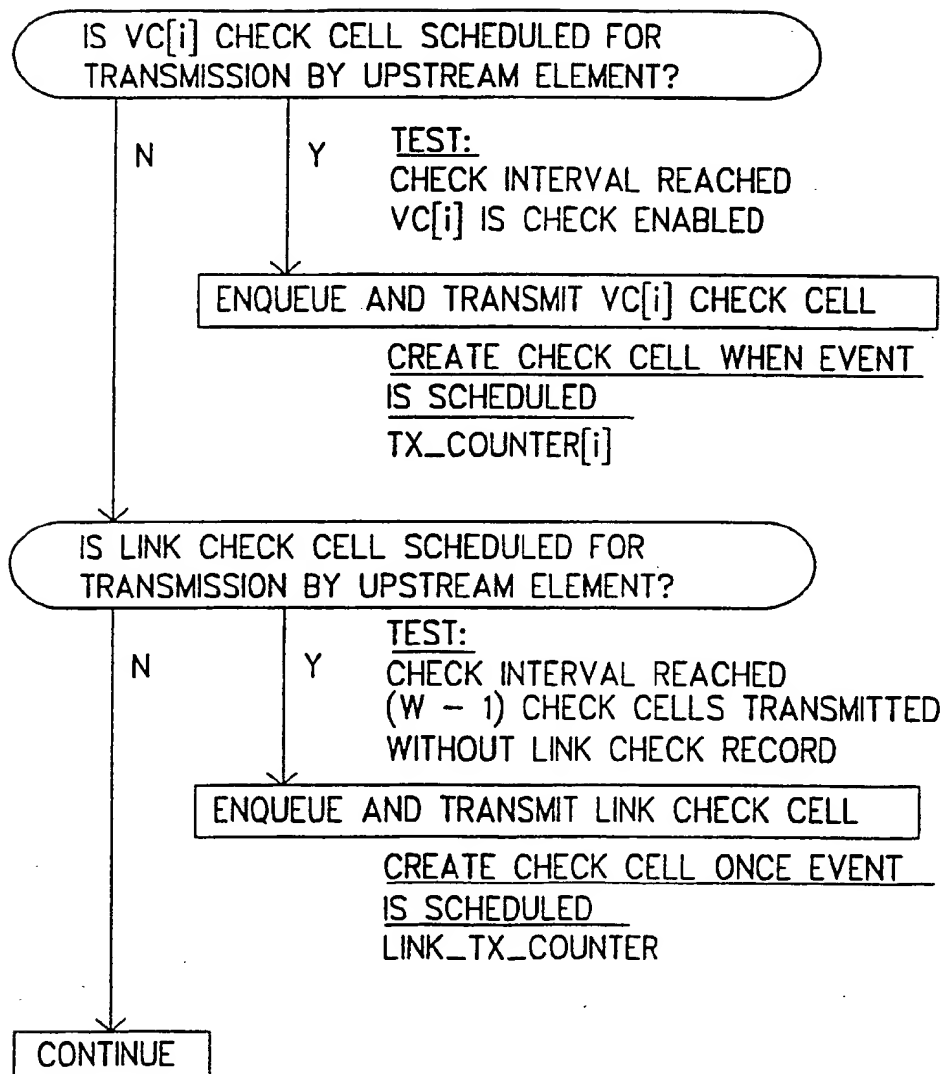


FIG. 9A

14/23

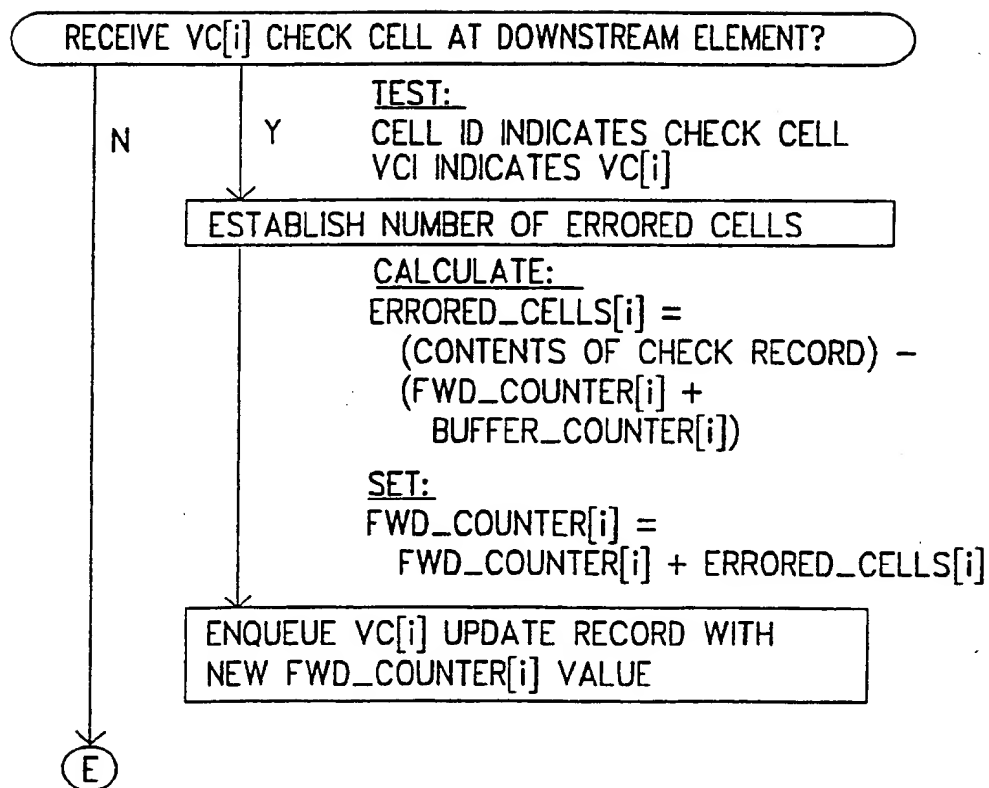


FIG. 9B

15/23

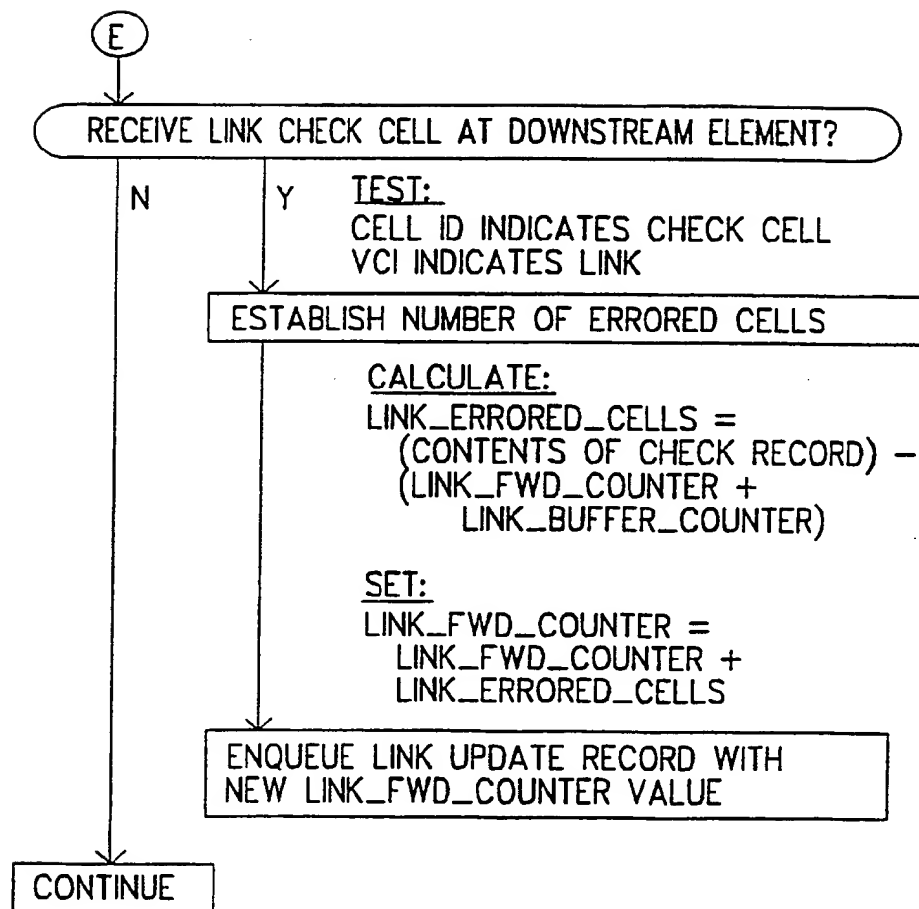


FIG. 9C

16/23




FIG. 10


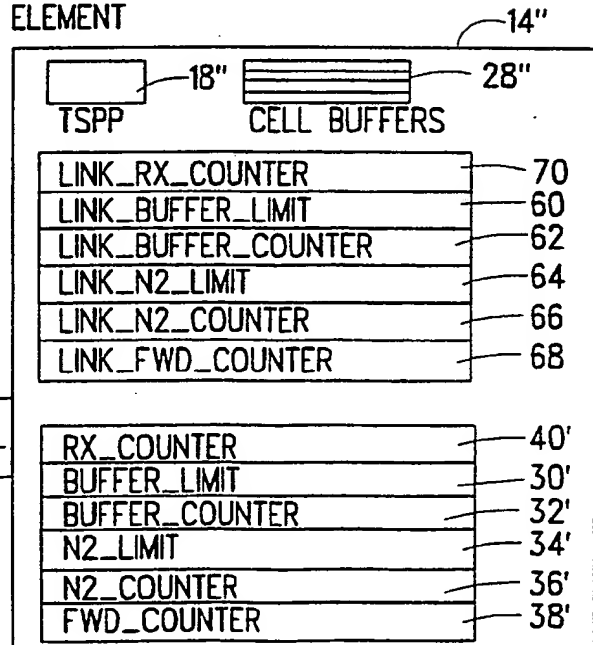
UPSTREAM
ELEMENTDOWNSTREAM
ELEMENT

FIG. 11

17/23

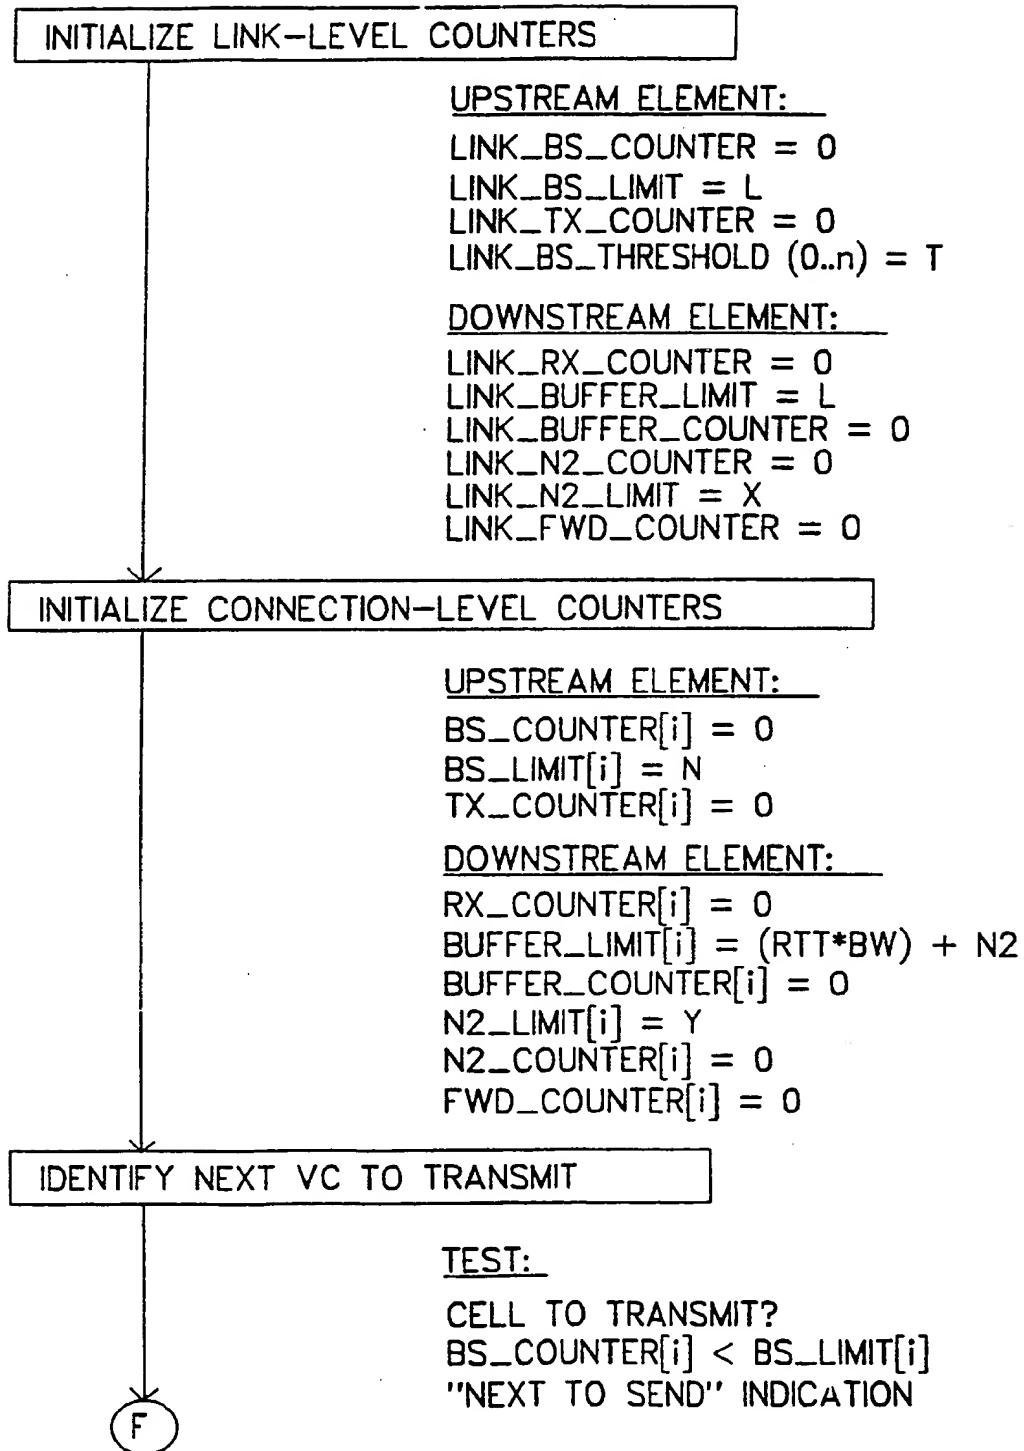


FIG. 12A

SUBSTITUTE SHEET (RULE 26)

18/23

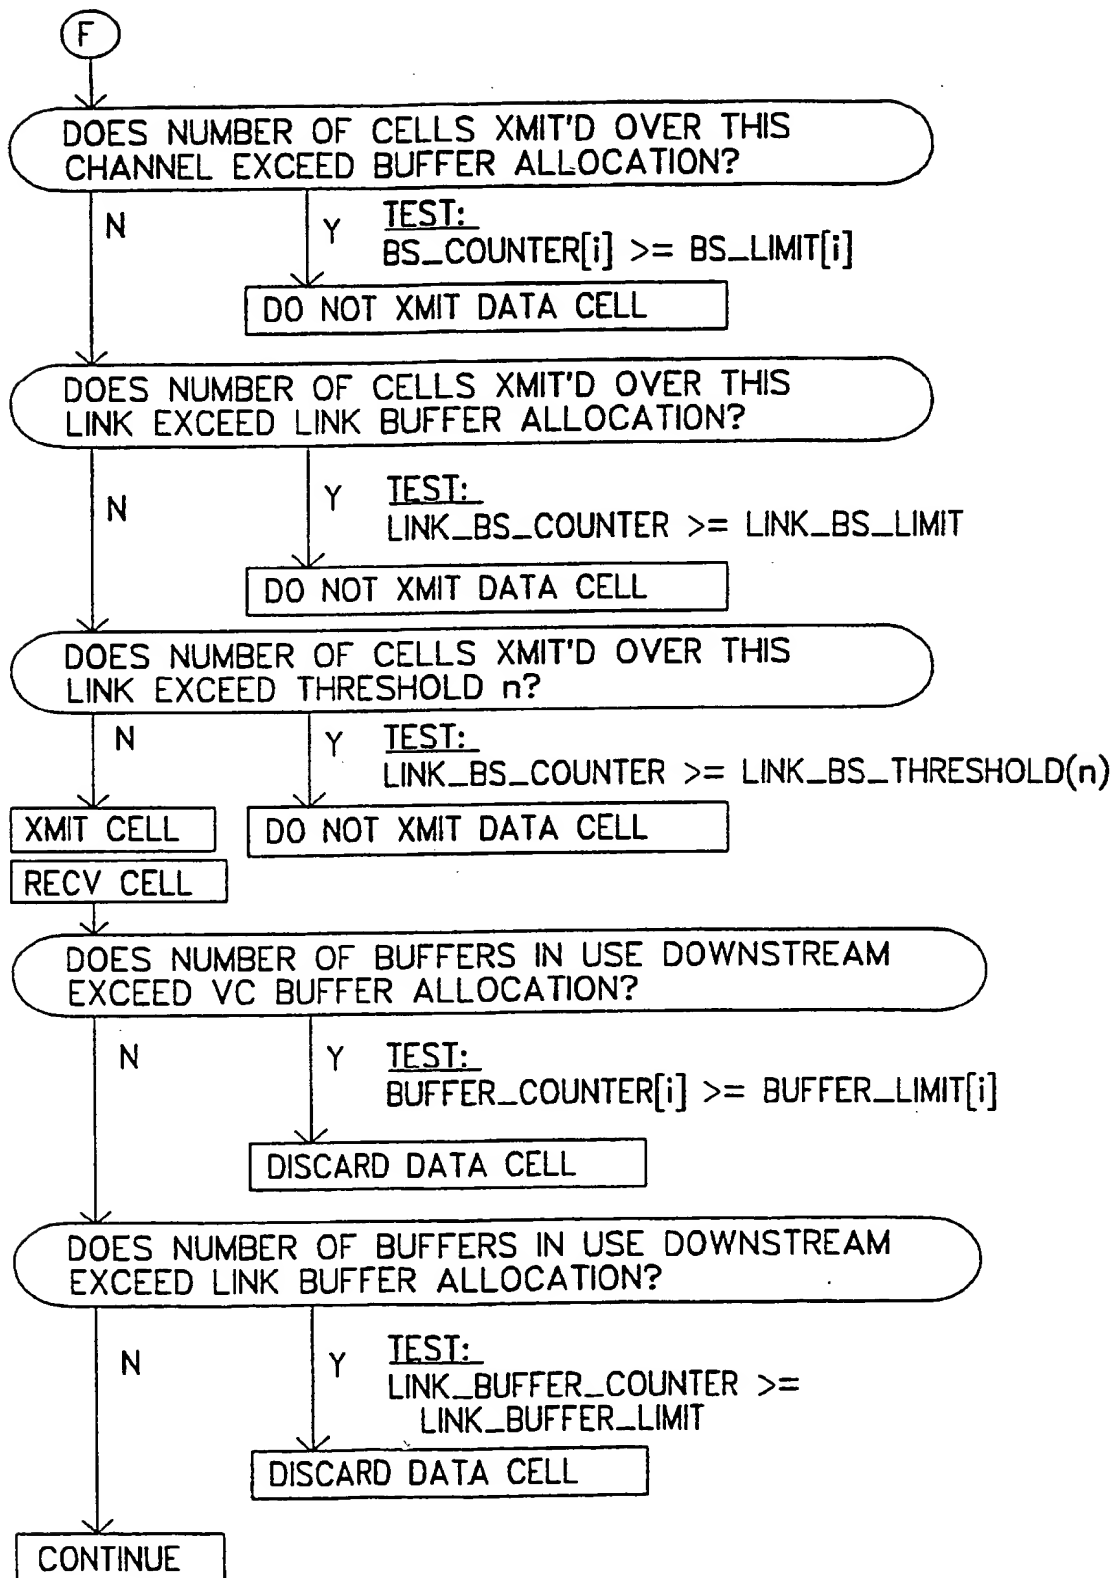


FIG. 12B

SUBSTITUTE SHEET (RULE 26)

21/23

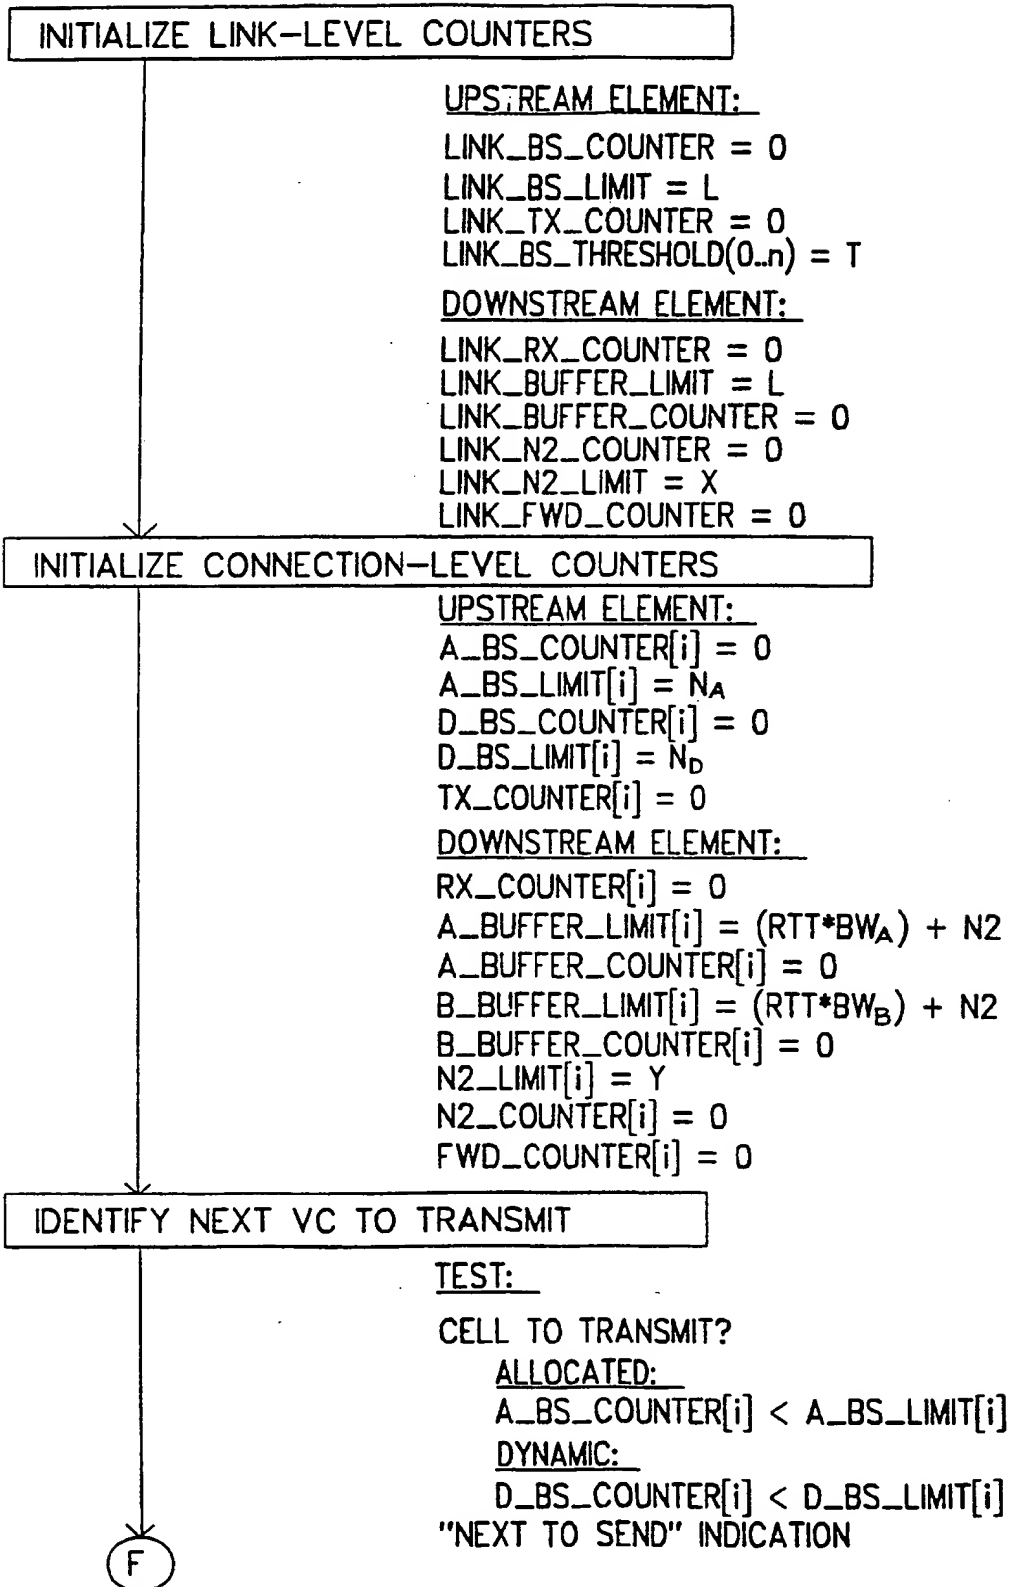


FIG. 15A

SUBSTITUTE SHEET (RULE 26)

22/23

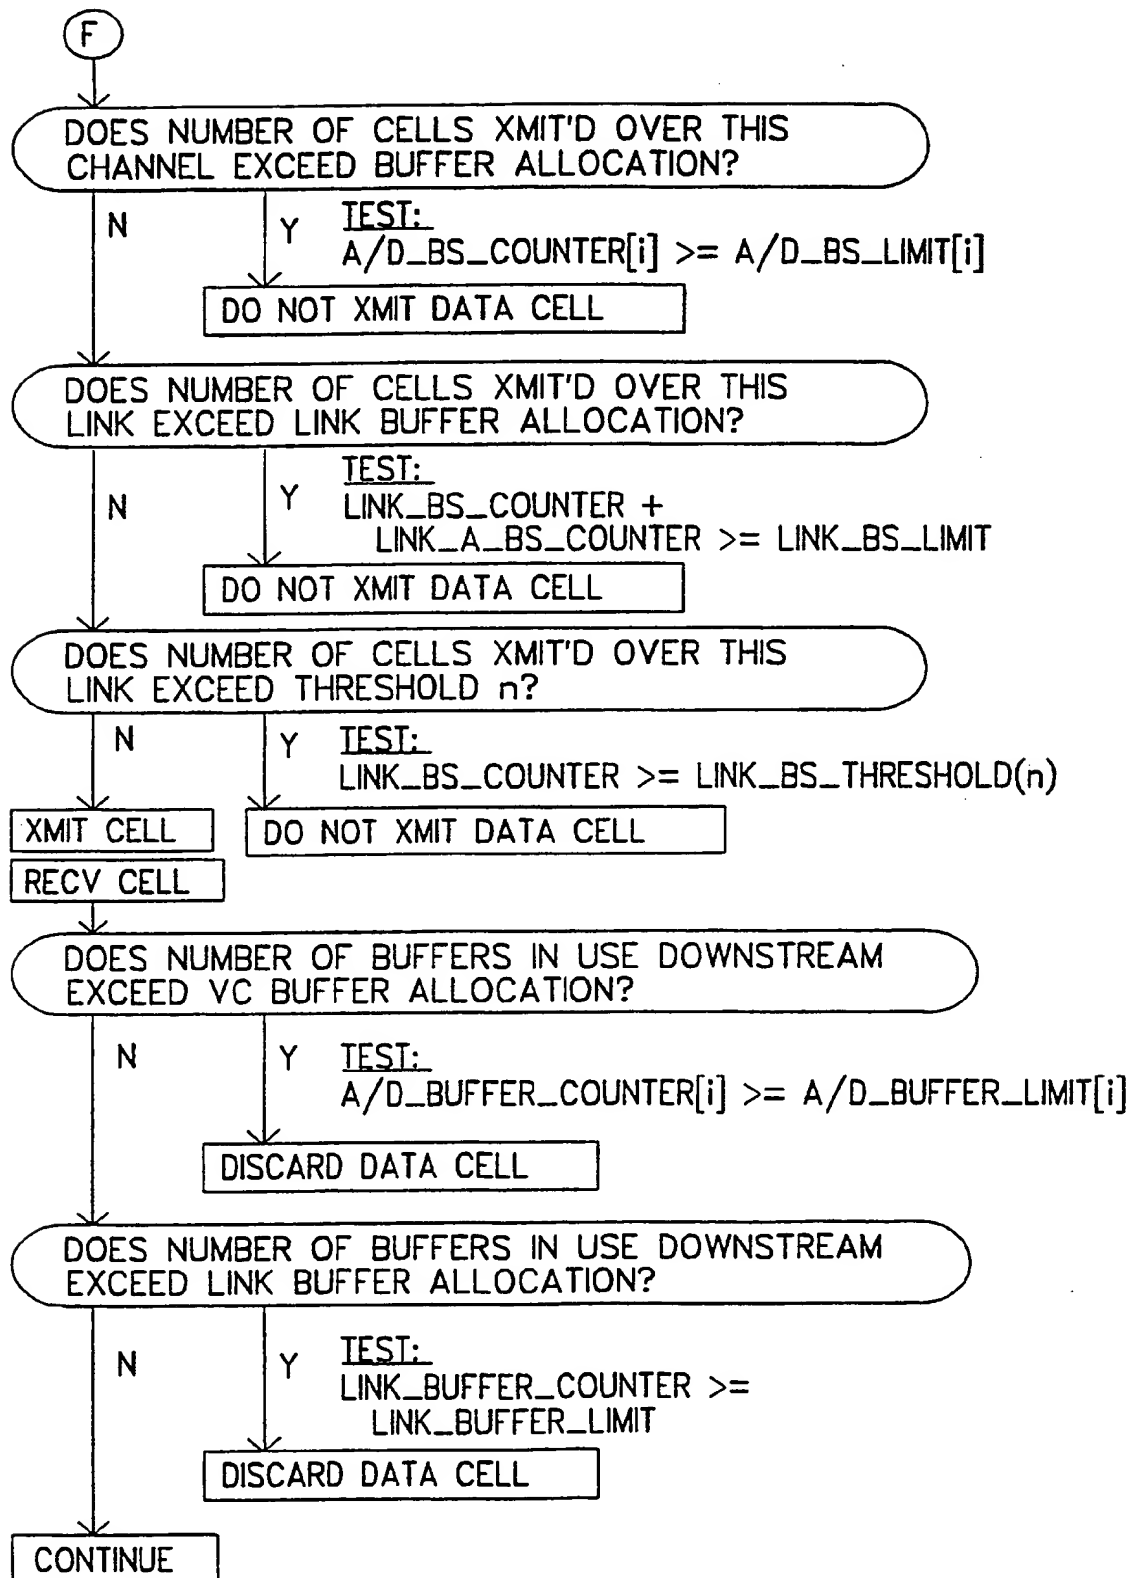
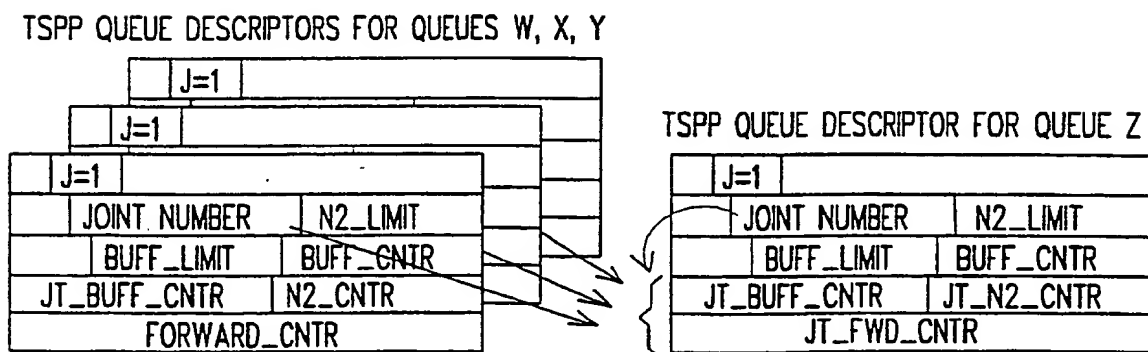



FIG. 15B

SUBSTITUTE SHEET (RULE 26)



PCT

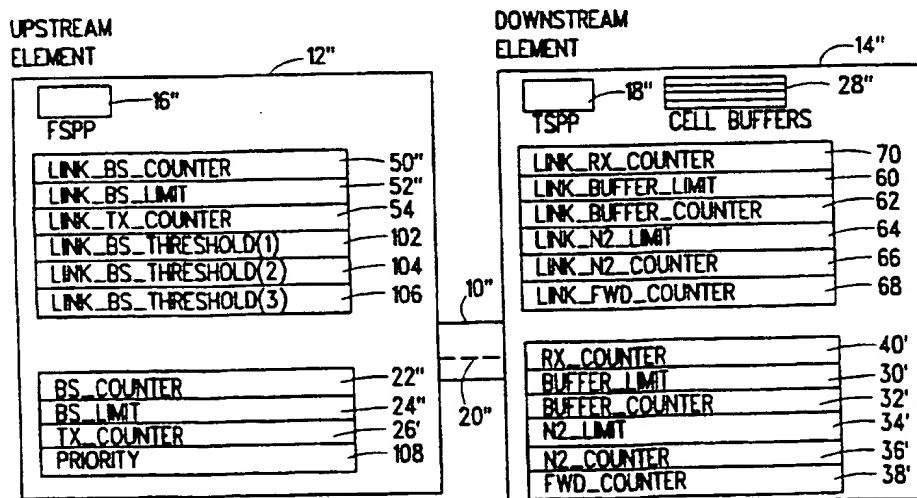
WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04L 12/56		A3	(11) International Publication Number: WO 97/03549
			(43) International Publication Date: 6 February 1997 (06.02.97)
(21) International Application Number: PCT/US96/11920		(81) Designated States: AL, AM, AT, AU, AZ, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TR, TT, UA, UG, US, UZ, VN, ARIPO patent (KE, LS, MW, SD, SZ, UG), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 18 July 1996 (18.07.96)			
(30) Priority Data: 60/001,498 19 July 1995 (19.07.95) US			
(71) Applicant (for all designated States except US): ASCOM NEXION INC. [US/US]; 289 Great Road, Acton, MA 01720 (US).			
(72) Inventors; and (75) Inventors/Applicants (for US only): HAUSER, Stephen, A. [US/US]; 106 Farms Drive, Burlington, MA 01803 (US). CALDARA, Stephen, A. [US/US]; 220 Horsepond Road, Sudbury, MA 01776 (US). MANNING, Thomas, A. [US/US]; 26 Summer Street, Northboro, MA 01532 (US).			
(74) Agents: LEBOVICI, Victor, B. et al.; Weingarten, Schurgin, Gagnebin & Hayes, Ten Post Office Square, Boston, MA 02109 (US).			
		Published With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.	
		(88) Date of publication of the international search report: 03 April 1997 (03.04.97)	

(54) Title: PRIORITIZED ACCESS TO SHARED BUFFERS



(57) Abstract

In a link-level flow controlled system, a method and apparatus providing the ability to partition a buffer resource among multiple prioritized buffer subsets (28) through definition of at least one threshold (102, 104, 106), the buffer resource being shared by a plurality of connections. Different category of service levels, in terms of delay bounds, are thus enabled. The presently disclosed link-level flow controlled system provides for zero cell loss. The shared buffer resource is divided among N-priority-pools, defined by N-1-threshold-levels, each priority-pool attributable to a respective category of service. Link-level counters and registers, disposed in a transmit element, as well as an indication of priority level associated with each connection, are employed in realizing the shared buffer resource.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US96/11920

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : H04L 12/56
US CL : 370/231, 235, 395
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 370/231, 235, 395, 229, 230, 232, 233, 234, 412

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- A	US, A, 4,727,537 (NICHOLS) 23 February 1988, abstract, col. 2, lines 14-38.	22 ----- 1-21, 23-32
X,P ----- A,P	US, A, 5,528,591 (LAUER) 18 June 1996, abstract.	22 ----- 1-21, 23-32
X,P ----- A,P	US, A, 5,515,359 (ZHENG) 07 May 1996, abstract.	22 ----- 1-21, 23-32
A	US, A, 5,083,269 (SYOBATAKE ET AL.) 21 January 1992, see entire document.	1-32

☒ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	*X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*G* document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means	
J document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

17 JANUARY 1997

Date of mailing of the international search report

25.02.1997

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

MIN JUNG

Telephone No. (703) 305-4363

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US96/11920

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A,P	US, A, 5,533,020 (BYRN ET AL.) 02 July 1996, see entire document.	1-32
A,E	US, A, 5,541,912 (CHOUDHURY ET AL.) 30 July 1996, see entire document.	1-32

Form PCT/ISA/210 (continuation of second sheet)(July 1992)*